



City Research Online

City St George's, University of London

Citation: Kwan Tai Yeong, E. (2004). The performance of technical analysts and technical forecasting. (Unpublished Doctoral thesis, City, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/30658/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

THE PERFORMANCE OF TECHNICAL ANALYSTS
AND TECHNICAL FORECASTING

STEVEN KWAN TAI YEONG

submitted in fulfilment for the degree of
Doctor of Philosophy

Faculty of Finance
Cass Business School
106 Bunhill Row
London EC1Y 8QZ

April 2004

ACKNOWLEDGMENTS

I wish to thank Professor Roy Batchelor for his help and guidance throughout the production of this thesis. Without him I would never have been able to start this project and without his suggestions, corrections and advice I would never have finished this work. He has given me a valuable opportunity to return to Cass Business School again after graduating in 1989, allowing me to broaden my knowledge base further.

TABLE OF CONTENTS

Chapter 1:	Objectives of thesis	1
Chapter 2:	Introduction of technical analysis	5
	2.1 Introduction	5
	2.2 Subjective methods	5
	2.2.1 Support and resistance	6
	2.2.2 Trend and Trendline	7
	2.3 Mechanical trading rules	8
	2.3.1 Moving average	9
	2.3.2 Momentum indicators	10
Chapter 3:	Literature survey of technical analysis	13
	3.1 Introduction	13
	3.2 Literature broadly classified into four categories	13
	3.2.1 Usage of technical analysis	14
	3.2.2 Profitability of technical-type trading rules	17
	3.2.3 Performance of technical analysis products	34
	3.2.4 Value of technical analysts data	37
	3.3 Summary and conclusions	39
Chapter 4:	A weekly survey of technical analysts	43
	4.1 Introduction	43
	4.2 Technical analysis	44
	4.2.1 Technical analysis Tools	45
	4.2.2 Evidence on technical analysis	49
	4.3 Our survey	52
	4.4 Survey results: Accuracy, Rationality and Profitability	56
	4.4.1 Forecast accuracy	56
	4.4.2 Forecast rationality	63
	4.4.3 Profitability	65
	4.5 Use of technical indicators	72
	4.6 Conclusions	81
Chapter 5:	Daily foreign exchange commentaries versus trading rules	83
	5.1 Introduction	83
	5.2 Studies of technical analysis	84
	5.3 Data	89
	5.3.1 Daily Commentary	90
	5.3.2 Trading recommendations	94
	5.3.3 Spot Exchange Rates	95
	5.4 Forecast accuracy and calibration	99

5.5	Profitability of recommended and mechanical trades	105
5.6	The logic of recommended trades	115
5.7	Concluding comments	122
Chapter 6:	Judgmental bootstrapping of trading strategies in the bond market	124
6.1	Introduction	124
6.2	Data and experts	126
6.2.1	The robustness of profits	131
6.2.2	Tests of expertise	134
6.3	Technical indicators	138
6.3.1	Survey of methods used by analysts	139
6.3.2	Technical indicators	141
6.3.3	Candlestick charts	144
6.4	Order response models	146
6.5	Beyond the judgmental bootstrap	158
6.6	Conclusions	161
Chapter 7:	Summary	163
References		167

CHAPTER 1

OBJECTIVES OF THESIS

The aim of the thesis is to evaluate the ability of technical analysis to predict movements in financial markets. This study is of obvious practical interest in that technical analysis is used intensively by market practitioners. It is useful to know whether there is any objective evidence that it works. The study has also proved timely in that technical analysis has stimulated a small but insightful programme of academic research in the past decade, and our work adds to that line of research. This study differs from previous academic research in one important way. All of our empirical work derives from information on the forecasts and trading recommendations of analysts themselves. This contrasts with most earlier studies, which try to mimic the forecasts of technical analysts by applying mechanical trading rules. Specifically, we utilise data from (a) a specially conducted survey of a group of analysts through 1998, and (b) unique data sets on daily published forecasts and trading recommendations by a leading provider of technical commentary through the years 2000-2001. Our analysis of this data strongly suggests that technical analysis does have value, and that the behaviour of technical analysts cannot be modelled using simple (or even quite complex) mechanical trading rules.

A widely accepted definition is that "Technical Analysis is the study of market action, primarily through the use of charts, for the purpose of forecasting future price trends" (J.J.Murphy, 1986). In contrast to "Fundamental Analysis", technical buying and selling strategies are based on the observation of past history activities, extracting market psychology from price patterns. The topic is important because technical analysis is by far the most common method used for short term forecasting by traders in financial markets. In spite of this, until recently there has been very little serious academic work on the value of technical analysis.

The prime reason for the paucity of academic work is the fact that technical analysis does not involve well-defined statistical procedures. Rather, technical analysis is an umbrella term for a very diverse collection of techniques, some quantitative and some judgmental, most with little scientific basis, and often sold with exaggerated claims about their likely success. This has made technical analysis an easy target for ridicule, most notably in Burton Malkiel's classic *Random Walk down Wall Street* (Malkiel, 1973).

In recent years, several factors have caused researchers to take technical analysis more seriously. From a number of "forecasting competitions" it is now well understood that a catholic approach helps improve forecast accuracy in a wide variety of business applications. The poor performance of many popular linear time series forecasting methods, notably Box-Jenkins analysis, in these competitions has shown that overselling is not unique to the world of technical analysis. Nonlinear modeling and forecasting methods are now in vogue. Indeed, in the financial markets, it is now recognised that many different regimes can be at work across a single time series of market prices, so it looks much more reasonable to use a set of tools rather than search for a single underlying model.

The availability of time series data based on high-frequency financial market prices has made it easier to make objective assessments of competing forecasting methods, and a number of academic studies have exploited this to evaluate technical analysis. However, this work has focussed on the narrow area of easily-replicated mechanical trading rules, such as moving averages and filters (for example, Brock et. al. 1992), and a few well-defined turning point patterns, such as the "Head-and-Shoulders" (Osler and Chang, 1995). Just as it would be hard to argue that a single-equation regression analysis would provide a good test of the value of conventional econometric forecasting, we argue that this focus on quantitative rules does not adequately reflect the complexity of the way technical analysis is applied in practice.

In this thesis we instead analyse information provided by technical analysts themselves, rather than predictions generated by notional mechanical rules. This information comes from two sources.

First, we conducted our own weekly survey of a panel of analysts making exchange rate and stock market forecasts through 1998. The survey includes information on directional predictions, volatility (trading ranges), and critical support and resistance levels. The survey also asks respondents what techniques they have used, so that these can be correlated with the accuracy and profitability of the predictions.

Secondly, we have extracted the daily commentary on currency and bond markets provided by the technical analysis team of a major international data vendor and information company. These comments are disseminated throughout the international foreign exchange and bond markets. They provide an opportunity to test the operational value of support and resistance “chart points” using hourly intra-day data, give us a better understanding of how these levels are used by FX analysts as compared to their trading recommendations.

The trading ideas drawn from the bond trading recommendations are matched by ordered response models by a subset of technical indicators. While some academic research exists for currency and stock markets, to our knowledge this is the first systematic investigation of the performance of technical analysis in the bond markets.

The plan of the thesis is as follows. In Chapter 2 we introduce and discuss the main techniques used by technical analysts. In Chapter 3 we survey the limited academic literature on the value of these rules. In Chapter 4 we describe the survey that we conducted among a panel of analysts, and assess the results. In Chapter 5 we describe the daily foreign exchange market commentary, and test the value of published support and resistance levels. In Chapter 6 we examine the bond trading patterns with order response models and to see whether models can mimic the judgment processes

of analysts. Chapter 7 summarises our findings, and discusses possible extensions to our study.

CHAPTER 2

INTRODUCTION OF TECHNICAL ANALYSIS

2.1 INTRODUCTION

There are many charting methods which fall under the category of technical analysis. In this section, I focus on the more traditional types used by analysts as suggested in the weekly survey done in 1998, chapter 4 in this thesis. I separate the topic into two categories: 1) Techniques based on subjective interpretation of the market and there are two popular types – a) Identifying support and resistance levels (requested in our survey) and b) the art of drawing trendline. 2) Mechanical trading rules that can be easily programmed into a computer; we focus on a) Moving averages and b) Momentum indicators. As the names suggested, these are used to measure rate of change of time series of market prices.

2.2 SUBJECTIVE METHODS

Most technical analysis techniques fall under this category. Prior to the advent of computers, pure observation of time series of market prices was the most convenient way of detecting trading sentiment. Technical analysts believe that anything that can possibly affect the market price of a market instrument is actually reflected in the price of that instrument. It follows, therefore, the study of price action is all that is required (J.J. Murphy 1986). Another two premises on which technical approach is based are 1) prices move in trends and 2) history repeats itself. The whole purpose of charting the price action is to identify trends in the early stages. As the study of market action has to do with the study of human psychology, a similar pattern appears on the price chart could have the same consequence as before.


2.2.1 SUPPORT AND RESISTANCE

Prices in the market fluctuate due to imbalances in supply and demand. Generally speaking, a *support* is a price level where buying is strong enough to interrupt or reverse a downtrend. This is an area where buying optimism will outweigh selling pressure. Trading activity consistently bounces off a particular price level where demand outweighs supply.




Supportive

A resistance is a price level where selling is strong enough, to interrupt or reverse an uptrend. It is an area, where selling pessimism will outweigh buying pressure. Trading activity is hindered repeatedly at a particular price level on the upside as further demand pressure fails to materialise.



Resistance

Once a support is violated, its role could potentially change to offer resistance on subsequent pullbacks in the trend. The vice versa applies for a resistance level.



To decide the “strength” of support and resistance, there are no rules of telling us how to differentiate levels but suffice to say, the more technical reasons the better the level is. Below are some examples of technical indicators which could potentially “strengthen” the support or resistance level. There are:


- Psychological level e.g. a round number
- Previous highs and lows
- Moving averages
- Trendline
- Congestion

The purposes are to detect trend change or the continuation of the underlying trend. The reason for the former is selling and buying pressure is expected to accelerate once a perceived key support/resistance is broken. A firm breakout could suggest a trend change although this is a subjective decision as to what constitutes a firm break. To detect the continuation of the underlying trend, a firm reaction from a previous support or resistance would confirm the importance of the direction. If it is a correction, i.e. the overall direction of the dominant trend is still intact, prices should be quickly arrested by key chart levels acting as support or resistance whatever the case might be to allow the underlying trend to continue at the prevailing direction.

2.2.2 TREND AND TRENDLINE

A trend is the prevalent direction of the market in a given time frame. However, markets do not simply trade in a straight line, interrupted with pauses after aggressive moves to form a succession of peaks and troughs but, this doesn't mean the dominance trend is weakening. It is the direction of these peaks and troughs that constitute the trend i.e. successive higher peaks and troughs would constitute an up trend and the reverse for a downtrend.

There are three directions of a trend: up, down and sideways with the latter corresponding to a flat trading range.



In an uptrend, a trendline is drawn by extending a rising line drawn from the two previous key support levels. The second of these obviously has to be higher than the previous trough, in order to establish the presence of an uptrend and this line is then extended into the future. In order to validate the existence of a trendline, it is preferred to be tested on a third occasion with a rebound needed to confirm but this is seldom the case in real live situation as the market might accelerate with momentum once a trend is established. In an uptrend, this is known as a support line. The same is true for confirming a trendline in a downtrend. In this case it is known as a resistance line.

The concept of a trend is that a trend in motion is more likely to continue than to reverse. That is to say once the presence of a trend has been established it will tend to exist until a trend reversal signal is appeared. A trendline can help to identify the extremities of correction. The assumption here being that a correction will find support in an uptrend and resistance in a downtrend at the trendline and the break of a trendline could suggest the first signal of a trend change. Therefore, this can be used to establish good buy/sell areas as long as it is not broken or some traders might prefer to use a “filter” for the breakout to get further confirmation of the breakout.

2.3 MECHANICAL TRADING RULES

With the advent of computers, measuring the rate of change of prices can be easily done. Trading rules can be easily tested for profitability with the access of high

frequency data . This in fact is one of the main reasons why most mechanical trend following systems are in use today.

2.3.1 MOVING AVERAGE

The simple definition of a moving average is the sum of n data points (i.e. prices) divided by n:

$$(x_1 + x_2 + x_3 + \dots x_n)/n$$

With each new data point the first is dropped to calculate an average for the last n points. This indicator can then be placed on the same price chart to compare current price action with the average of the past n periods.

The application of moving averages can be broadly divided into three categories:

1) Support and Resistance – The moving average displays a smoothed version of the price series and therefore cuts out volatile movement to concentrate on the underlying direction. In a downtrend, the moving average will lie above the price to provide resistance, and in an uptrend, this could provide support. The main purpose is to see whether the current price is able to maintain above or below the average of the n period. 2) Crossover – To use moving averages as part of the trading strategy, the crossover of the average by price should be watched. For example in a downtrend, as prices cross above the average, this gives a signal of a change in sentiment to potentially bullish. The criteria for a valid crossover can be a close above the average or a full session trading above the average. The opposite is true for a potential bearish signal when prices drop below the average line. 3) Parameters - The choice of parameter for the average is important for it to give valuable signals. A very short period average would retain the “choppiness” as prices. At the same time too large a parameter will generate late signals. Therefore a correct balance needs to be found and usually require trial and error in the process of optimising.

Same as the idea of prices crossing over a moving average line, the implication is the same for cross over between two moving averages with different time parameters. A two moving averages system uses a long term moving average and a short term moving average. In an uptrend the longer period average is below the short term moving average to indicate increasing rate of ascent and vice versa in a downtrend. The potentially bearish trading signal is given when in an uptrend the short-term average crosses below long term moving average. On the other hand, the positive signal in a downtrend is when the short term average crosses above the long term average to indicate changing market tone from bearish to potentially bullish. As for a three moving averages system, this consists of a long term, medium term and short term. In an uptrend, the initial signal of a trend change to potentially negative is given when the short-term average crosses below the medium term average. This may be used as an opportunity to liquidate a previous long position or stay out of the market until a stronger trend reversal signal is evident. The major negative signal would be when both the medium term and short term cross below the long-term average to confirm a change in the underlying trend. The same rules applied in a bearish trend reversal.

2.3.2 MOMENTUM INDICATORS

Momentum is used to describe the underlying strength of a market. This is to estimate sentiment of the market crowd, and changes in this can give clues as to future price behavior.

There are many ways of measuring momentum, but these can be broadly classified into two types. The first, often overlooked but still very important, is by simple inspection of the price chart itself. Momentum change can manifest itself in various ways, the most obvious of which is change in slope of the chart, but can also include such features as reversal days (higher high plus lower close in an uptrend) signaling internal weakness. The other major category comprises a large selection of indicators,

the Relative Strength Index by W. Wilders and Stochastics by George Lane are the most popular in the marketplace.

The Relative Strength Index (RSI) measures the relative internal strength of a price series.

The formula for constructing the RSI is as follows:

$$RSI = 100 - \frac{100}{1 + RS}$$

$$RS = \frac{\text{average of } x \text{ number of days up closes}}{\text{average of } x \text{ number of days down closes}}$$

By construction the RSI will have the upper boundary at 100 and lower at 0. A positive slope signals bullish sentiment and a negative slope suggests bearish. However, the indicator will rarely if ever reach these extreme values. In normal market conditions it is common to use the reading of 70 as overbought and 30 as oversold. The most important point to note about the overbought/oversold concept is that this should only be used as a warning rather than a confirmed trend reversal signal. In a strong bullish uptrend, oscillators could stay “artificially” overbought for a long period of time. Therefore the first and foremost indicator is price action itself.

In trading range scenarios indicators such as the RSI can be particularly useful although overbought and oversold signals should be confirmed by charting analysis on the price chart e.g. breakout of trendlines. At times, divergence occurs when the oscillator line and the price line diverge from one another and start to move in opposite directions. This indicates the change in prices is not confirmed by the change in momentum. Positive divergence has higher lows in the indicators but lower lows in prices implying sentiment is changing from bearish to bullish. Negative divergence has higher highs in prices but lower highs in the indicator.

One of the most popular momentum studies favoured by technical analysts is George Lane's Stochastics. Two lines are used in this process - the %K line and the %D line.

The formula for 5-day stochastic is as follows:

$$\%K = 100 [(C-L5)/(H5-L5)]$$

where C is the latest close, L5 is the lowest low for the last five days, and H5 is the highest high for the same five days.

$$\%D = 100 \times (H3/L3)$$

where H3 is the 3 day sum of (C-L5) and L3 is the 3 day sum of (H5-L5), i.e. a very short term moving average. Note: %K indicates the level of the most recent closing price in relation to the range of the last five periods in the past.

Due to the construction of the formula, the momentum oscillator fluctuates between 0 and 100, The higher the closing value in relation to the trading range (rather than the absolute value), the higher the %K line. Same idea as moving averages technique, a crossover from the %K line above/below %D line would signal sentiment change and the subsequent turning from the %D would confirm this directional move. Divergence again has the same implications as in the RSI. Positive divergence has higher lows in oscillators but lower lows in prices, signalling weakness in the dominant downtrend and vice versa for a negative divergence.

The technical tools mentioned above are popular in the technical analysis but this is only a small subset of the techniques available to traders. A more comprehensive study of various technical techniques is advised to look at Technical Analysis of the futures markets by John J. Murphy (1986)

CHAPTER 3

LITERATURE SURVEY ON TECHNICAL ANALYSIS

3.1 INTRODUCTION

In this literature survey, we look at the past research done on the subject of technical analysis. Previous investigations can be separated into four categories. Some studies simply report the extent of the usage of technical analysis, without commenting on its merits. Some try to go further and look at the profitability of rules available to technical analysts but not necessary the way in which these are actually applied in a real trading environment. Very few investigate the performance of financial products using technical analysis and negligible research done on the value of technical information provided by technical analysts, this is the void that this thesis is trying to shed light on.

3.2 LITERATURE BROADLY CLASSIFIED INTO FOUR CATEGORIES

The survey starts by looking at the research on the usage of technical analysis. This approach is to look at the survey of technical analysts. Secondly, I look at the ones that concentrate more on the technical technique. For example moving averages, patterns recognition like head and shoulders, the five point patterns that include various formations like channels, wedges, diamonds, triangles, double/triple bottoms. Also, the investigation of "market efficiency" with filter rules is also included in this section. Thirdly, the focus is on a few papers on the performance of financial products using technical analysis and finally, we look at the value of data provided by technical analysts themselves.

In all cases, the results are by no means consistent and further investigation is warranted.

3.2.1 USAGE OF TECHNICAL ANALYSIS

The prevalence of technical analysis in short term forecasting is well documented in Taylor and Allen (1992), based on a survey conducted in London in 1988, on behalf of the Bank of England, among chief foreign Exchange dealers.

They mailed a questionnaire to 402 named institutions and a total of 213 questionnaires were returned. Deleting the 16 which were turned back by the Post Office and the 33 were returned marked as being inapplicable, this gives a response rate of 60 per cent, enough to draw some reliable conclusions about the use of technical analysis.

The responses confirmed what everyday observation of dealer behaviour would suggest - that 90 per cent of respondents place some weight on this form of non-fundamental analysis when forming their view of exchange rate expectation at the horizon of intraday to one week. At the longer forecast horizon, of one to three months, or six months to one year, the focus shifted to the fundamentals.

The majority held chartist and fundamentalist approaches to be complementary to a greater or lesser degree, only 8 per cent being seen these as being mutually exclusive.

However, technical analysis is seen as much more decision-relevant. For firms employing both economists and technical analysts in-house, the advisory role is 61.5% used economists and 34% to help position-taking. In contrast, technical analysts are used by only 38.5% as advisors, and 45.3% to guide position-taking.

These findings are confirmed in the paper by Lui and Mole (1998), based on a questionnaire survey made in February 1995 among FX dealers in Hong Kong, again concerning the usage of technical and fundamental analysis for their FX forecast.

From the original list of 812 members, a total of 153 fully completed questionnaire were returned with a response rate of 19% (presumably the higher response rate in the Taylor and Allen study reflected the greater status of the Bank of England). The findings reveal that more than 85% of respondents use both Technical and Fundamental analysis, and skew towards Technical Analysis at the shorter horizon from intraday to one month. It is also revealed that Technical analysis is felt to be slightly more useful in forecasting trends than fundamental analysis but significantly more useful in predicting turning point. The latter registered an average score of 7.26 (with 0 for not important/useless and 10 for very important) and a standard deviation of 1.69 of responses indicates that respondents have the smallest diversity of opinion with regard to this superiority. The corresponding figures for the usefulness of Fundamental analysis are 4.63 with 2.24 s.d.

Their major findings may be summarised as follows:

- 1) At all time horizons, a very high proportion of respondents place some weight on both fundamental and technical analyses when forming views. At shorter horizons, there exists a skew towards reliance on technical as opposed to fundamental analysis, but the skew becomes steadily reversed as the length of horizon considered is increased.
- 2) Dealers perceived value in using both fundamental and technical analysis to predict both trend and turning. Technical analysis is particularly strong in forecasting turning point.
- 3) In the use of technical analysis in forecasting trends and turning points, the most common length of historical period used by dealers is 12 months and the most used data frequency is daily data.
- 4) Interest rate news is found to be relatively important fundamental factor, while moving average and/or other trend following systems are mostly used technical techniques. Nevertheless, they are both given less weight than news about central bank intervention in influencing intraday exchange rate movements.

In the discussion paper by Thomas Gehrig and Lukas Menkhoff (2003), the survey done on FX dealers and fund managers has highlighted technical analysis has gained importance over time, it has by far the greatest importance in FX dealing and is second in fund management. This survey suggests the importance of technical analysis used in the shorter term forecasting and in line with previous surveys, fundamentals are more dominated in longer term horizon but, the interesting point here for the shortest term view, flows are cited to be the key focus for market participants. This finding is in line with the project done by Yin-Wong Cheung, Menzie Chinn and Ian Marsh (1999) on UK-based FX dealers. When asked about to select the single most important factors that determine exchange rate movements in the intraday horizon, the three popular factors mentioned were bandwagon effects, over-reaction to news and speculative forces. The survey shows very little evidence of systematic differences of opinion between technical and fundamental analysis.

In J. Frankel and K. Froot (1990), the authors offered a different insight into the usage of technical analysis by looking at the question of "overshooting" of the US dollar rate. The most dramatic aspect is the period from June 1984 to February 1985 where the dollar surged 20% over this interval as other observable factors that are suggested in standard macroeconomic models were all moving in the different direction.

One view is that the appreciation may have been an example of a speculative bubble - that is not determined by fundamentals but rather the outcome of self-confirming market expectations. There is evidence that investors have heterogeneous expectations, survey of the forecasts of participants in the foreign market show wide dispersion at any point in time. It has long been remarked that if there are traders who tend to forecast by extrapolating recent trends then their actions can exacerbate swings in the exchange rate. Technical analysts are thought to use rules that are extrapolative such as buy when 1 week moving average crosses above the 12 weeks moving average.

The model of speculative bubbles could be due to the shifted weight away from fundamentalists and towards chartists according to the authors. Euromoney magazine

runs a yearly August review of between 10 and 27 foreign exchange forecasting firms. This reveals that in 1978, 18 forecasting firms described themselves as relying exclusively on economic fundamentals and only 2 on technical analysis. By 1985, the position has been reversed only 1 firm reported of relying on fundamentals and 12 on technical analysis. It may indeed that the shift focus of forecasting techniques is the source of changes in the demand of the dollar.

The objective of the studies above was simply to measure the extent of the use of technical analysis. Many other studies have tried to assess whether technical analysis methods have any value to its users.

3.2.2 PROFITABILITY OF TECHNICAL-TYPE TRADING RULES

Given that many chart techniques are subjective and it is difficult if not impossible to model the technical analysts' behavior fully. Some of the methods used by analysts are not easily compute or mechanical enough for testing - examples are Elliott Wave and Pattern recognition.

However, several attempts to evaluate technical trading rules themselves have been made over the past years, mostly selecting the "branch" of technical analysis perceived to be less subjective and these are usually categorised as trading rules like Moving Averages and Filter Rule or Trading Range Break. The former is the average closing price of the exchange rate over a given number of previous trading days and the idea is to generate trading signal when the short moving average crosses the longer moving average. For example, when the short term 5 day moving average crosses from below (above) of the 20 day moving average would generate a buy (sell) signal, the lengths of the moving averages are chosen by the technical analysts. As for the filter rules, buying (selling) an instrument when it rises (falls) x percent above (below) its previous local maximum (minimum). In trading range break, buy and sell

signals are given when prices moved away from a defined range. There are attempts to test the validity of the technical pattern of a Head and Shoulders as well.

When investigating the subject of "Random Walk", the basic hypothesis of the theory is that successive price changes in individual securities are independent random variables. That is past history of a series of changes cannot be used to predict future changes in any "meaningful" way. That is to say if investors want to know whether history of prices can be used to increase expected gain, in a random walk market, no mechanical trading rule, applied to an individual security would consistently outperform a policy of simply buying and holding the security.

Since some of the methods used in the subject in Technical Analysis are mechanical e.g. the moving average rules, it is not surprising to conclude that if the market movement is random, technical analysis should be fruitless in predicting future movement. Another way of looking at Technical Analysis would be that since almost all techniques are based on historical data, the whole "philosophy" behind the thinking is "not appropriate".

In a classic paper, Alexander(1961,1964) uses a filter technique on movements in stock prices is based on the idea that if the daily closing price of a particular security moves up at least x per cent, buy and hold the security until its price moves down at least x per cent from the subsequent high, at which time sell and go short. The short position is maintained until the daily closing price rises at least x per cent above a subsequent low at which time one covers and buys. Alexander formulated the filter technique to test the belief, widely held among market professionals that prices adjust gradually to new information.

In his earlier articles, the tests were done for the filters ranging in size from 5 to 50 percent. The tests covered different time period from 1897 to 1959 and involving closing prices for two indexes, the Dow Jones Industrials from 1897 to 1929 and Standard and Poor's Industrial from 1929 to 1959. In general, filters of all different

sizes and for all the different time periods yield substantial profits – indeed profits significantly greater than those of the simply buy and hold. This led Alexander to conclude that the independence assumption of the random walk model was not upheld by his data.

In Fama and Blume (1966), further examination of the work by Alexander's filter technique was carried out. Alexander's previous filter rule results were first questioned by Mandelbrot (1966) about the overstatement of the profitability of the filters. Although he later adjusted the assumption of trade could always buy at a price exactly equal to the low plus x per cent and sell at the high minus x per cent, the results are still difficult to interpret. The main reason is the difficulty of using price indexes to adjust for the effect of dividends.

Applying Alexander's filter technique to the individual securities of the Dow Jones Industrial Index between the date from January, 1956, to April, 1958, but are usually about the end of 1957, to Sep 26 1962, there are altogether thirty samples with 1200/1700 observations per samples in this research. The results are in general inferior to the technique of filter rules (twenty-four different filters ranging from 0.5 per cent to 50 per cent). For example when commissions are taken into account, only four securities have positive average returns per filter. Obviously, when commissions are omitted, the returns from the filter technique are improved greatly but still not as large as the returns from the simply buy and hold strategy.

The above is inconsistent with the result of Alexander finding with superior performance of filter rules compared to buy and hold when commissions are omitted due to the difficulty of adjusting indices with dividend. Under the buy and hold policy, the total profit is the price change for the time period plus any dividends that have been paid. Dividends simply increase the profitability of holding shares. Under the filter technique, however, the investors alternate long and short positions. In a short sale, the borrower of the securities has to reimburse the lender for any dividends that are paid while the short position is outstanding reducing the profit of short sale.

The results presented in this paper show that filter rules only surpass the buy and hold policy for dividend adjusted for only two securities. The breakdown of the returns before commissions for the long and short transactions add further evidence that the simple filter rule probably cannot be used to increase expected profits. The filter rules are disastrous for the investors for short positions and long position average returns are smaller than the buy and hold policy as well.

When commissions are included, all filters below 12% and above 25% produce negative average returns. Taking filters within 12% and 25% the average is small compared to the buy and hold policy. These results support the conclusion that the filter techniques cannot be used to increase the expected profits of the investors who must pay the usual brokerage commissions.

But subsequently in 1988, in the paper by Sweeney (1988), the inferior results by Fama and Blume are further investigated. This paper claim that studies of the 1960s tended to understate filter rule returns relative to buy and hold.

The test selects a subset of the Fama and Blume stocks that look more promising in their work and follow these stocks from 1970 through 1982. For this sample of stocks, it appears that significant profits can be made by investors with low but feasible transaction costs. In particular floor traders who avoid the use of specialists can achieve these profits while those who pay even the lowest commercial rates very likely cannot; borderline significant profits may exist for institutional money manager.

The paper's rule considers only long position while Fama and Blume have the investor short a particular security whenever he is not long in it. The short positions perform poorly, so avoiding them raises the measured returns and saves transaction costs. Fama and Blume looked at all 30 Dow-Jones Industrial stocks for the late 1950s and early 1960s, they averaged filter profits across stocks. The approach of this paper instead looks at only the winners in one period and asks whether there is persistence so they remain winners in later periods. The individual stocks that looked

like winners in the investigation of Fama and Blume study for the $\frac{1}{2}$ of 1 per cent rule are re-examined for the period from 1970-1982. All gave filter rule returns that were statistically significantly better than buy and hold when floor trader transaction costs of $\frac{1}{20}$ of 1 per cent are used for each one-way transaction. For an equal weighted portfolio, the filter significantly beats buy and hold even for transaction costs as high as $\frac{3}{20}$ of 1 per cent.

But the above results are sensitive to transaction costs and to whether closing price is an unbiased estimate of the price at which one can buy or sell. An interesting issue of why floor traders are still able to make substantial profit is due of the cost of the seat in the exchange; the cost of the seat is the present value of the profits that could be made. Rules are only implemented by market participants to the extent justifiable on risk grounds. Another mentioned here is that attempts to capture the profits will eliminate them before the investor can execute the transactions.

In the paper by Levich and Thomas (1991), the filter rules were further re-examined. This time using the futures contracts for the period of 1976 – 1990 of the currency pairs of British Pound, Canadian dollar, German mark, Japanese yen and Swiss franc. A bootstrap method of generating numerous samples size n from the original data set to test the null hypothesis of no information and signals in the original sequence was also being looked at.

The quotations are on closing settlement prices from the International Monetary market of the Chicago Mercantile Exchange. A single time series is then assembled by bring together quotations on successive near term contracts with filter rules of size of 0.5%, 1%, 2%, 3%, 4% and 5% and three moving average cross-over rules of 1day/5 day, 5 day/20 day and 1day/200 day.

The profits associated with the generation of buy and sell signals using filters and moving averages are significant. Over the entire 15-year sample period, every size filter results in positive profits for every currency. Average profit in the Canadian

dollar across all filters is 2.0%, substantially less than the average for other currencies where results range between 6.9% and 8.1%. The results are much the same for the moving average rules which led to average profits of 2.7% for the Canadian Dollar and between 7.0% and 9.0% for other currencies. In terms of transaction cost, the likely cost of transacting in the currency futures market is about 2.5 basis points (0.025%) per transaction for a large institution, a more conservative estimate would be roughly 4.0 basis points. Since the volume of trading is considerably smaller apart from the short moving average rule and the small filters, the transaction cost do not, according to the authors, alter the results much.

The rank of the filter rules profits for the actual series in comparison to the randomly generated series (10000) for each currency pairs in this case has given striking results as well. In 19 of the cases, the profits of the actual series rank in the top 1% (9900 and above) of all the simulated series. In 6 further cases, the rank is in the top 5% (9500-9899). The remaining five cases rank lower, but in no case lower than the top 21% of the simulated series. Thus, in 25 of the 30 cases, the authors reject the hypothesis that there is no information in the original series that can be exploited for profit by the filter rules. The results are much the same for the moving average rules as well.

These results strongly suggest that the actual exchange rate series contained significant departures from serial independence that allowed technical trading rules to be profitable. If the actual series has been generated randomly, the simulations suggest that average profits would be close to zero, same as the average profit of the 30 randomly selected cases.

Possible explanations for the persistence of trading profits are the presence of central bank intervention that tends to retard exchange rate movements only temporary and the profitability of trend following rule may be the result of excessive speculation that cause prices to follow.

Moving away from research on filter rules, Brock, Lakonishok and LeBaron (1992) used the Dow Jones Index from 1897 to 1986, the test of the moving average and trading range break were presented. The findings were quite encouraging for the moving average buy signals generating returns not likely to be performed by the null models like random walk, AR (1), GARCH-M and EGARCH. There were five moving average combinations chosen according to the authors, the "popular ones" and they were 1-50 (short term 1 day and longer term 50 days), 1-150, 5-150, 1-200 and 2-200.

The tests were done in several different ways. The first rule was called the variable length moving averages (VMA). This initiates buy (sell) signals when the short moving average crosses above (below) the longer term moving average. This rule is further modified by the introduction of a band around the moving average i.e. long and short positions are initiated when amount of the crossover is larger than the band. The second rule focuses on the crossing of the moving averages and maintained position for a specific period of time. This method was called the fixed length moving average (FMA). In this exercise, a ten day period was selected i.e. returns during the next ten days are recorded and other signals during these ten days are ignored.

Looking at the buy signals, for the VMA rules, the returns, for the crossover and crossover with the band, are all positive with the average one-day return of 0.042, approx. 12 per cent at an annual rate. For the second moving average test of FMA, all the buy signals generated positive returns as well.

For the sell signals, all the returns are "negative" for the VMA method for the crossover and same for the FMA, sells are all negative.

Apart from testing moving averages, the method of trading range breakout (TRB) was investigated as well. Buy (sell) signals are generated when the price level moves above (below) local maximum (minimum). Local maximums and minimums are computed over the preceding 50, 150 and 200 days. Like the moving average

techniques above, a 1 percent band was introduced and also the 10 day holding period. The buy return is positive across all rules.

In the paper by Sullivan, Timmermann and White (1997), further research of the Brock, Lakonishok and LeBaron (1992) paper is carried out. There are two important elements added in this investigation. First, the introduction of the White's Reality Check bootstrap methodology to evaluate simple technical trading rules while quantifying the data-snooping bias and fully adjust for its effect in the context of the full universe from which the trading rules were drawn. Secondly, the price data of the DJIA was extended for another 10 years, from 1987 – 1996, the Brock et al investigation was from 1897 – 1986.

The problems of Data-Snooping can be seen from various angles. 1) When given set of data (trading rules) is used more than once for the purposes of inference statistics or model selections, there is always the possibility of getting a satisfying results simply due to chance rather than any merit inherent in the methods. 2) Data snooping need not be the consequence of a particular researcher's effort, it could be due to survivorship bias i.e. as time progresses, rules that performed well receive more attention. 3) If enough trading rules are considered overtime, some rules are bound to work by pure luck even though they do not genuinely possess predictive power.

By extending for another 10 years of the original data used by Brock et al, this would allow the investigation into the question of whether the previous profitable 26 rules identified by Brock et al would be profitable in the out of sample test.

Sullivan et al address this issue of data snooping by constructing a universe of nearly 8000 parameterizations of trading rules. Also using the same set of data Brock et al used to investigate the potential data snooping in their experiment.

Selecting the best performing trading rule based on the mean return criterion, the best trading rule of the BBL's 26 trading rules universe stands up to closer inspection of

data snooping effects. For the data period from 1897 to 1996, the best trading rule was a 50day variable moving average rule with a 0.01 band. This has outperformed the buy and hold strategy. But, from the full universal rule by Sullivan et al, an even better rule is found of a standard 5-day moving average. However, the superior of the best performing rule is not repeated in the out of sample experiment covered the 10 year period from 1987 – 1996. This is also true when applied on the S&P 500 futures from the period of January 1984 – December 1996.

Two explanations cited for this “inability” to perform when using out of sample data. First, this could be due to the stock market crash in October 1987. According to the paper, this argument cannot be rejected outright but it is never the less of a rather long sample of 3291 days. Also, a large swing could improve performance as short position would be initiated.

Another reason suggested was the market has become more efficient during the out of sample period and hence opportunities have disappeared. This scenario matches up with the increase of liquidity in the stock market.

Ramazan Gencay (1996) investigates moving average rules with a band i.e. a fixed % of deviation, between the short and the long averages. The paper uses the daily Dow Jones Industrial Average Index from January 1963 to June 1988 to examine the linear and non-linear predictability of stock market returns with buy and sell signals generated from the moving averages. The results indicate the adoption of the band eliminates noisy buy and sell signal and improve the quality of the out of sample forecast. The results indicate that non-linear conditional mean specifications of the past buy/sell signals of the moving average rule with a band provide forecast improvements for the current returns over the linear model which uses past returns as regressors.

The paper by Pruitt and White (1988) takes a different approach by combining moving averages with other technical signals. The research uses the trading system

approach i.e. a combination of mechanical tools. CRISMA (Cumulative volume, Relative Strength, Moving Average) identifies security "target" by using the three most commonly employed technical filters which are based upon a stock's relative strength compared to the S&P500, cumulative volume, and 50 day and 200 day moving averages of prices. It attempts to measure and "triple confirm" upward momentum. Buy recommendations are issued on potential securities only when the fourth and final penetration filter has been satisfied.

This is how the trading model works for a buying recommendation. First, the 50-day price moving average graph must intersect the 200-day price moving average from below when slope of the latter graph is greater or equal to Zero. Second, the Relative Strength graph from the beginning to ending point over the previous four weeks, must have a slope greater than or equal to zero. Finally, the cumulative volume graph from the beginning to ending point over the previous four weeks must have slope greater than zero. If these three criteria are satisfied, stock is purchased when its price reaches 110% of the level intersection of the 50 and 200 day moving average graph and exit when prices decline below the 200 day moving average or rise above 120% of the established level.

During the period from 1976 to 1985, a total of 204 stocks hit these criteria in the University of Chicago's CRSP daily data types. The 204 stocks recommended by the CRISMA trading system were assumed held for a total 4970 security-days. The mean and median length of security holding were 24.4 days and 18 days, respectively. The longest holding period was 113 days and the shortest was two days. The four return generating models used here are the mean-adjusted return model, the market-adjusted returns model, the ordinary least squares market model and the Scholes-Williams market models.

Using the 0 transaction cost with various model return, the mean daily excess 0.1066% (26.65% annualised) to 0.1426% (35.65% annualised) and with 2% transaction cost, the annualised figures were between 6.13% and 15.13%. The results

suggest CRISMA ability to outperforming a simple buy and hold strategy over a significant period of time even adjusting for timing, risk and transaction costs.

Lukac, Brorsen and Irwin (1988) look at the scenario that the concentration of the usage of certain types of technical analysis in the form of mechanical system by traders effects the futures prices. Markets from time to time are being subjected to unprecedented waves of one way buying or selling.

After conversations with traders, 12 trading systems were selected for testing. The results suggested trading systems trade on the same day significantly more often than would randomly be expected but the actual percentage of trades that occur on the same day is small. Furthermore all the systems are on the same side of the market, significantly more than would be randomly expected, and the returns are significantly positively correlated. If the users of these systems are numerous, they have the potential to move the market prices, but this would only happen for a minority of the trades.

Moving average is again the technique being used in the research by Silber (1994). Although the research in this paper is less to do with the validity about the subject of Technical Analysis, it concentrates on the hypothesis that technical trading rules work in markets where there is price smoothing behaviour by non-profit-maximising participants like central banks intervention. Fears of central bank activity couple with occasional massive expenditures by central banks might introduce sluggish price adjustments that could be exploited by traders.

The idea is that if technical analysis works in the foreign exchange markets because of the government intervention, then it might also work in trading short term interest rates such as Federal Fund, Treasury bills and Eurodollar time deposit. Moreover, if that were the only reason for the success of technical trading, technical trading rules should not work in markets without significant government intervention like gold, equities. This study tests these conjectures by simulating the profitability of the

moving average trading rule as applied to futures contracts on foreign currencies, short term interest rates and other actively traded commodities.

The moving average rule used in this research is the moving average crossover i.e. buy and sell signals are based on the relationship between a short term and a long term moving average of prices. The rule require a long position when the short term moving average crosses the long term one and short position when the shorter average falls below the long average. The selecting process of the moving average used is first select the most profitable average for year 1 and then use the optimal combination for year 2. The next step is to re-optimize to select the most profitable of the year1 and year2 combined and use this for year 3, this sequence is continued through the end of the sample period. Due to the unlimited amount of combination, the authors restrict the research to short term of 1,2,3...15 and long term of 16,18,20,22....200.

The instruments used for the research are futures contracts and there are separated into two groups, one is perceived to have the influence from the "non-profit maximising" behaviour and the other has free movement. The first group consists of German Mark, Swiss franc, Japanese Yen, British pound, Canadian dollar from the Chicago Mercantile Exchange, also Eurodollar from CME and the three month Sterling from the London International Financial Futures Exchange. The second group has gold and silver from Comex, US T-bond from Chicago Board of Trade crude oil from the New York Mercantile exchange and the S&P500 from CME. The authors do recognise that there might contain elements of price smoothing of the T-bond due to intervention at the short end of the yield curve and the OPEC for the oil. The number of years of data ranged from 8 to 12 within the period from 1979 to 1991, apart from the crude oil of only 7 years.

The results showed profitable returns for the first group but not the second even taking into account of the adjustment of the transaction cost, the latter is assumed bid-ask is equal 1 tick for each contract except crude oil and gold for 2 ticks. Positive

returns from the moving average are also recorded a much higher returns than the buy-hold strategy for all contracts apart from the three-month Eurodollar. Four out of the seven contracts that associated with price smoothing are significantly positive to 5%. Sharpe ratios are also supportive to the hypothesis that returns per risk are related to the existence of price smoothing behaviour in the market. One important caveat mentioned by the authors is that sticking to trading positions that may have little support from fundamentals require dedication and discipline.

The central bank intervention was investigated by Andrew Szakmary and Ike Mathur (1997). In this paper, moving average trading rules are utilised in both futures and spot foreign currency markets to show that significant positive profits can be earned in four of five currencies examined during the period of June 2 1977 to June 28 1991. Regression results demonstrate that central bank intervention is strongly associated with profitability of trading returns. Christopher Neely and Paul Weller (1999) highlighted this on technical analysis and central bank intervention. It shows that technical trading rules can make use of information about the US FX intervention to improve out of sample for two of four exchange rates examined. Rules tend to take positions contrary to official interventions and unusually profitable on days prior to intervention.

In Levy (1971), a different approach was used in analysing the “effectiveness” of Technical Analysis. This is the approach where all the possible combinations of a five point reversal patterns are identified by a computer algorithm of the testing series and then to test the results after the breakout. In this exercise, there are altogether 548 New York Stock Exchange securities covered the time period from July 3 1964 through July 4 1969.

There are altogether 32 combinations of a five point patterns and the key is to identify the reversal points and then match the five points to one of the 32 formations. The “Reversal” is defined as the quantity of “ $a + bV$ ” where a and b are constant and V is the volatility of the individual stock as measured by the arithmetic average of the day-

to-day percentage price change over the most recent 131 day period. The breakout was defined as a price movement which penetrate the fourth reversal point of the pattern. For the 9383 patterns identified, performance was measured from the occurrence of the pattern breakout through twenty-six weeks into the future.

The results are not too encouraging for Technical Analysis as neither the best nor the worst of these thirty-two rules performed very differently from the market. Note: the rate of return relative to the market, where the "market" is gauged by the daily geometric average of the adjusted prices of the 548 stocks in the research file). After taking into trading costs into account, none of the thirty-two patterns showed any evidence of profitable forecasting ability in either bullish or bearish direction. More over, the most bullish results tended to be generated by those patterns which are classified as bearish in the standard textbooks on charting and vice versa.

However, there are a number of points which are needed to be taken into account as well. 1) The use of daily close rather than high-low range. 2) The specifications of the reversal rule. 3) The definition of "Breakout". 4) The failure to require a specified minimum percentage breakout prior to taking a long or short position.

In the staff reports published by Osler and Chang (1995), the researchers were concentrating on the Head and Shoulders pattern to evaluate its predictive power. The primary defining characteristic of a Head and Shoulders pattern is a sequence of three peaks with the highest in the middle as the "Head" and the left and right peaks are referred to as "Shoulders". The pattern is completed when the price path crosses the straight line connecting the two troughs separating the head from right to left, also known as the "neckline"

The test was done on daily spot rate of six currencies against the dollar: The yen German mark, Canadian dollar, Swiss franc, French Franc and pound, from period from March 19 1973 to June 13 1994.

The research methodology is first to identifying the Head and Shoulders pattern from the computer-implemented algorithm by tracing out the zigzag pattern in the data according to technical books. Then to define a peak as a local maximum, at least x percent higher than the preceding trough is needed and a trough as a local minimum at least x percent lower of the preceding peak. The value of x is called the “cutoff”. Ten cutoffs with different values are chosen i.e. to scan the data ten times to gather the total number of Head and Shoulders, but eliminate duplicate patterns.

Also, a number of other criteria are set as well for the selection i.e. the height of the Head must exceed the height of the left and right shoulders. A pattern should occur following a trend and rules for vertical asymmetries and horizontal asymmetries are defined to get the “right” shape of the pattern. And, the “time limit” for breakout is also being considered to avoid establishing trading positions prematurely.

Once a Head and Shoulders pattern is identified, a position is entered when the price line breaks the neckline. Also, rules of exit of the position for examples the stop loss, bounce and unwind position after a specific number of days, are also being considered.

For the mark and the Yen, head and Shoulders based profits derived from actual foreign exchange data are significantly greater than those derived from artificial data but, for the Canadian dollar, the Swiss Franc, the French Franc and the pound are not significantly different from those derived from artificial data. However, speculating in six currencies simultaneously shows aggregate meaningful profits.

However, the Head and Shoulders signal is not too promising when used in the stock markets. In Osler (1998), trading strategy based on the pattern of Head and Shoulders is being “treated” as noise trading. The identification of Head and Shoulders traders as noise traders is based on two empirical results. First, Head and Shoulders traders substantially increases aggregate trading volume, amount to one quarter of a day’s trading volume. Secondly, Head and Shoulders trading is not profitable.

The analysis is based on a computer program that identified Head and Shoulders pattern based on certain defined characteristics. The algorithm finds about 27 confirmed Head and Shoulders pattern per firm. The data consist of prices and volume from 100 firms selected at random from the entire Centre for Research on Securities Prices equities data set, from the period of July 2, 1962 to December 31, 1993. Altogether 528 firms, 100 were selected.

For the argument of Head and Shoulders is not profitable and hence one of the “prerequisite” of being termed as noise trading. The suggestion here is that since the pattern if used as a signal for speculative trading, the goal of which is to make profit, a lack of profitability is sufficient to show that such trading does not qualify as rational speculation, and thus does qualify as noise trading. The null hypothesis is that the Head and Shoulder patterns are meaningless noise and the null is tested using bootstrapping method and here the result is not profitable.

Andrew Lo, Harry Mamaysky and Jiang Wang (2000) propose a systematic and automatic approach to technical pattern recognition. The findings of this approach to a large number of US stocks from 1962 and 1996 shows several indicators do provide incremental information and may have some practical value.

A genetic programming approach was being used by Neely, Weller and Dittmar (1997). The idea here is using genetic program as a search procedure for identifying optimal trading rules and then applying on the out of sample data. Six currencies were selected (\$/DM, \$/JPY, \$/£, \$/SF, DM/JPY and £/SF) in a sample period from 1975 – 1980 and then examine their performance over the period of 1981-1995. The advantage of this approach is the ability to construct a true out-of-sample test of the significance of the excess returns earned by the trading rules.

For all the exchange rates, the authors used 1975-77 as the training period, 1978-80 as the selection period, and 1981-95 for validation. There are 100 rules for each

currency. The results of excess returns are somewhat variable across currencies but they are all positive. If the "median" portfolio rule is used i.e. long position used if 50% or more of the rules signalled long and a short position otherwise, This trading portfolio improved the average excess returns significantly. Despite the difference in the frequency of trading, the transaction cost has not affected the result too much.

Looking at the best 10 trading rules for the \$/DM rate, the rule that did the best during the examination period performed the worst in out of the sample testing but rule two to ten all produce excess returns. However, rules that are generated by the genetic program have a rather complex nested structure and might not be the ones that technical analysts commonly used.

To discover whether the returns to the trading rules could be interpreted as compensation for bearing systemic risk, the authors also calculated the betas with the four benchmarks of the MSCI world equity index, the S&P 500, the commerzbank index of German equity and the Nikkei. Note: Betas are the coefficients from regressing monthly excess returns from a portfolio of 100 rules for each currency on the monthly excess returns for each of the equity indices. The results suggest excess returns were not compensation for bearish systematic risk.

Also, the best rules derived from the \$/Dm rate were also very profitable for the other \$ pairs and the less well on the DM/JPY could suggest differences between dollar market and the DM/JPY market.

There is a source of information asymmetry, specific to the foreign exchange market that may play a role in generating profitable trading opportunities. Central bank intervention is one explanation although during the period of early 80s, this is not a frequent occurrence and banks have private information about future fundamentals in the form of changes in monetary policy. Another possibility would be some form of market inefficiency.

Christopher Neely and Paul Weller (1999) have highlighted the ability of trading rules to generate excess returns for three of four EMS exchange rate over the out-of-sample period 1986-1996 especially allowing the use of information about interest rate differential. For example, a DEM/ITL rule would take the form of "Take a long position if the minimum interest differential (Italian minus German) over the last 4 days exceeds 3.88". However, the results of the trading rules cannot be duplicated by commonly used moving average rules or filter rules.

Have traditional technical trading rule profits in the currency markets declined over time, this is the question asked by Dennis Olson (2003). The paper tests whether moving average trading rule profits have declined over the period from 1971 and 2000. Rules are optimized for successive 5-year in sample from 1971 to 1995 and tested over subsequent 5- year out-of sample periods. Results show that risk-adjusted trading rule profits have declined over time from the average of 3% in the late 1970s and early 1980s to about zero in the 1990s.

3.2.3 PERFORMANCE OF TECHNICAL ANALYSIS PRODUCTS

An early empirical study is Murphy (1986), which looks into the trading records of futures funds managed by fund managers using technical analysis. These are funds managed by professional technicians which invest in futures contracts and provide investors with the opportunity to take diversified positions in the futures market. The data for this study consist of 60 monthly observations from May 1980 through April 1985 on all purely technical futures funds listed in the first "Fund Review" section of *Commodities Magazine* (11 altogether).

The basic framework for testing and evaluating mutual fund performance is based on the work developed by Sharpe (1966), Hakansson (1971) and Jensen (1968). The idea behind Sharpe is that investors generally risk averse, so that the ratio of reward per unit of risk (volatility) is relevant. Hakansson advocated the geometric mean as a

useful measure of performance and Jensen has taken into consideration the usefulness of the fund in a portfolio context, and suggests evaluating performance using the abnormal return over the CAPM-implied return.

As a baseline, the author includes the naïve strategy of a portfolio of equal dollar long positions in the “second nearest” futures contract for each of 30 different types of commodities and financials listed on the US Exchanges. For the market proxy, two alternative portfolios are employed, with the first proxy being a simple 100% investment in corporate equities. The second proxy consisting of a portfolio of 60% equities, 30% corporate bonds and 10% T-Bond.

With Sharpe ratio being used to rank the fund, most funds earned positive returns but only one outperformed the S&P 500, implying concentrating of wealth in a futures fund is inferior to portfolio concentrating in a group of stocks. As for the Hakansson’s geometric mean, only five funds outperformed a risk free strategy of rolling over one month T-bills, and on average, funds did not do as well as the riskless benchmark. The Jensen index was not impressive as well, with no significant abnormal return found after added to a portfolio of stocks over the sample period. Also, there is no statistical evidence that any of the futures funds would improve the performance of a diversified portfolio of stocks and bonds.

In terms of evaluating the relative merits of technical analysis versus a naïve buy and hold strategy, the average technical fund on all three performance criteria were outperformed by the naïve strategy. But, only one technical fund was found to generate returns significantly superior to those of passive strategy.

On a gross basis, the results of technical funds are slightly more encouraging, generating statistically significant abnormal returns and also outperforming the naïve strategy, the S&P 500 and the T-bill. But this can be seen as only “sufficient” to at least cover brokerage and management fees.

In terms of analysing technical market commentaries, Brown, Goetzmann and Kumar (1998) investigates the track records of one of the Wall Street most famous chartist W. Hamilton.

This paper has a different concluding remark than the previous research done by Alfred Cowles (1944) on W. Hamilton. In the previous work of Alfred Cowles (1933), he provides strong evidence against the Hamilton's forecasting ability, his analysis is a landmark in the development of empirical evidence about the informational efficiency of the market. However, this paper concluded that the Dow theory applied by Hamilton over the period of 1902-1929 did yield positive risk adjusted returns. The difference in the results is apparently due to the lack of adjusting for risk. Cowles compares the returns obtained from Hamilton's market timing strategy to a benchmark of a fully inverted stock portfolio. In fact Hamilton's portfolio is frequently out of the market. Adjustment for systematic risk appears to vindicate Hamilton as a market timer.

The Dow theory states that the market movements may be decomposed into primary, secondary, and tertiary trends, the most important of which is the primary trend. This can be identified by a few key signals. First, a "trend" must be confirmed by both the industrial and the transportation indices. Second, extended movements sideways, called "lines", presage the emergence of a definite trend.

Apart from the adjustment of the systematic risk mentioned above, Cowles, performing of the non-parametric analysis of bullish forecasts and bearish forecasts of Hamilton's recommendation also lack of consideration of the efficacy of repeated bull forecast in a rising market and bearish forecast in a falling market. That is to say any sequence of positive or negative calls confirmed by a rising or decreasing would be reduced to a single datum.

The authors decode the 255 Hamilton editorials during the period from 1902-1929 as bullish (54% of the time), neutral (24%) and bearish (22%). To address the basic

question of Hamilton's timing skill, the authors examine how often the Dow beats the riskless rate over the interval following an editorial, conditional upon a bull or a bear call. The interval following the editorial is defined by the day following the editorial to the day of the next Hamilton editorial. The results are encouraging with calls right 110 times and wrong 74 times, especially in bear calls, showing strong evidence of association between Hamilton calls and subsequent market performance. From statistics, the Fisher's test is statistically significant at the 1 per cent level and the Henriksson Merton (HM) test gives compelling evidence that Hamilton was particularly effective in bear markets with the proportion of correct bear calls is much higher than those anticipated by chance.

3.2.4 VALUE OF TECHNICAL ANALYSTS DATA

In Allen and Taylor (1990), a survey was carried out over the period of Jun 1988 – Mar 1989 (38 weeks) for a one week and four weeks ahead forecast. A panel of analysts was telephoned (approx. twenty chartists participated) every week and their expectation with respect to the \$/£, DM/\$ and Yen/\$ recorded. Three points were made from the inspection of the survey figures. Prediction errors are noticeable greater at the four week horizon than one week. There is a tendency of forecasts to miss turning points. And there is a broad tendency to underpredict in a rising market and to overpredict in a falling market.

These features suggest that chartists' expectations are extrapolative. The example given was the DM/\$ rise in qualitative accuracy at the one week horizon between September to October exactly matches the establishment of a downtrend of this rate. But, the subsequent reversal of the \$ was matched by a fall in qualitative accuracy.

In measuring the accuracy among the Panel, the non-parametric test procedure which allows for matched samples shows reasonable evidence of systematic differences in forecasting performances across the panel. Only one chartist, named "M" in this paper consistently outperformed the median and also the Random walk model. The

median itself had a lower root mean square error than the majority of individual chartists, perhaps suggesting the consensus view is likely to outperform most individual views in aggregate although it fails to outperform the Random walk. On the other hand, forecasts based on time series models - ARIMA models and VAR systems - failed to outperform the forecasting panel.

Curcio and Goodhart (1992) examine the daily data on support and resistance levels and range predicted by chartists on Reuters. The test was done on three currency rates: Deutsche Mark, Sterling and Yen, against the dollar over the period from April 1989 to 29 June 1989.

One common agreement among technicians is that once a support or resistance level has been broken, this is a sign that a trend in that direction has started and that is likely to continue. From the authors' point of view, this branch of technical analysis testing would hopefully avoid the criticism of using arbitrary rules among a potential infinite set and from a technical analysts standpoint there is not the testing models unfamiliar in the technical circle.

Six trading rules were created and the idea is the same i.e. initiate new positions when prices move beyond the range. The position is then kept open until either the exchange rate moves back into the range or the range is revised to include the actual value of the exchange rate. The six trading rules are 1) The value of the support and resistance on the screen as the first boundary. 2) An outer 0.1% band is implemented to 1). 3) Used the forecast trading range. 4) 0.1% is applied to 3). 5) The band consists of 1) and 3). 6) The band consists of 2) and 4).

The results for this exercise are that the average returns from following both the buy and sell signals are always positive and higher than the average returns for the whole series. The average returns from following both signals are substantially higher than the whole series as well. Only buy signals noticed in the DM/\$ and JPY/\$ are profitable and only sell signals profitable for £/\$, this is suggestive of the ability to

generate good trading signals for the overall \$ positive trend. Also, using the moving average rules used in Brock et al (1992), profitable results were seen as well.

The value of support and resistance levels is further examined by the Carol Osler (2000). In this paper, the idea is to test the ability to predict intraday trend interruption when these are levels are approached/reached. Using data of three currency pairs, DM/\$, Yen/\$ and \$/£, from six active market participants of the period from Jan 1996 to Mar 1998, the finding is encouraging. The results indicate that intraday exchange rate trends were interrupted at the published levels subsequently more often than would have occurred had the levels been arbitrarily chosen. Despite their overall success, none of the firms correctly assessed the relative likelihood of trend interruption at the different levels and firms do not agree extensively with each other on the relevant signals.

3.3 SUMMARY AND CONCLUSIONS

The extensive use of technical analysis in the financial market is confirmed by the research done by Taylor and Allen (1992) and Lui and Mole (1998) on participants in the FX markets, in major financial centers in London and Hong Kong respectively. Relatively greater weight is given to the use of technical analysis for forecasting short term horizons, ranging from intraday to weekly (monthly also mentioned). But for the longer term outlook (over one month), participants are still biased towards fundamental analysis. Lukac, Brorsen and Irwin (1988) show that a technical analysis mechanical system, if implemented by many participants, has the potential to move market prices. The use of technical analysis was cited as a potential reason for the overshooting of the DM/\$ during the period of 1981-1985 noted by J Frankel and K.Froot (1990).

The superiority of technical analysis of predicting turning point cited in Lui and Mole is however not being totally agreed by other studies. The research carried out by

Brown et al (1998) on the market commentaries of chartist W.Hamilton has produced encouraging results of his ability as a good market "timer". However, this is not consistent with the early findings by a survey research done by Allen and Taylor (1990) of which cited technical forecasting as extrapolative and this again being mentioned in Frankel et al (1990).

When analyzing the profitability of using technical analysis, a noticeable shift of techniques and sophistication was seen over the past 10 years. This could be due to the advance of computer and software capability. However, the results are by no means conclusive.

In filter rules, the original promising results by Alexander (1961,1964) on stock indexes of the Dow Jones and S&P industrial were subsequently questioned by Mandelbort (1966) and being further investigated by Fama and Blume (1966) and then re-examine again by Sweeney (1988). In Fama and Blume (1966) the filter rule techniques were applied on individual securities of the Dow Jones Industrial Index instead, overcoming the difficulty of adjusting of dividend payment as in indices by Alexander, and the results were in general inferior. Sweeney took a slightly different approach by analyzing the "winners" from Fama and Blume and re-examined this in a different time period. This has produced consistent profitable results.

Using filter rules and moving averages as the technical trading rules was the methodology used in Levich and Thomas (1991) on currency futures. Here the profits associated with the generation of trading signals are significant. In Brock, Lakonishok and LeBaron (1992), the rules of moving averages were further investigated on a long data sample of the Dow Jones Index. Altogether 24 rules are tested including 9 rules from the method of trading range breakout. The results are positive. This however has provoked the question of data snooping and in Sullivan, Timmermann and White (1997), the rules from Brock et al were further examined for the data snooping bias. The results of Brock et al appear to be robust to data snooping however and indeed,

there are trading rules which performed even better than the ones considered by Brock et al.

In terms of knowing a bit more about how profitable trading rules perform from out of sample data, a genetic programming approach was used in Neely, Weller and Dittmar (1997) on spot currency rates. The results were somewhat variable but there are all positive. This is different to the inferior result from Sullivan et al when applying the best trading rule observed to out of sample data.

Moving away from the mechanical trading rules, the investigation of technical patterns and information provided by technical analysts are also seen. In Levy (1971) of testing a "five point" pattern of New York Stock Exchange securities, the results were poor but, acknowledged that the variation of the rules are enormous. In Osler and Chang (1995), a more specific computer algorithm was used to identify the pattern of Head and Shoulders in spot currency. The results are mixed for the currency pair investigated but, speculating in all the currency pairs (six pairs altogether) shows aggregate meaningful profits. The "profitability" of the Head and Shoulders pattern is however poor in the research done on company shares in Osler (1998).

Looking at data and information provide by technical analysts, research studies done on this type of reports are few and far between. Pretty promising results were cited by Curcio and Goodhart (1992) of range breakout rules using support and resistance levels offer by technical analysts. While investing money on futures fund managed by technical analysts, the results were less apparent by a different number of measurements done by Murphy (1986).

Overall, however, it is unclear whether technical analysis is useful. The subject has indeed been investigated by academics more rigorously over the past decade, attempting to focus on the more subjective and judgmental methods rather than simple mechanical rules. But, very often, the outcomes of the research seem to differ,

depending more on the type of market data used than on precise technique under investigation. It also appears that many tests are unrealistic. Many of the trading rules applied in academic studies continue to be artificial, and it is unclear whether real trades do take place in the market based on these rules. In addition, there are many possible trading rules, and it is not obvious which are important in practice. Applications of the rules are very often over-simplified by researchers, with none looking at combinations of rules. Similarly, although some studies do try to mimic the trading environment - e.g. taking into consideration of commissions and spreads, the role of money management in trading is generally ignored.

The lessons we draw from this survey are that we need to establish (a) which of the many possible technical tools are really used in the market, and (b) how exactly these rules relate to trading positions taken by analysts. These form the subject of Chapters 4 and 5 below.

CHAPTER 4

A WEEKLY SURVEY OF TECHNICAL ANALYSTS

4.1 INTRODUCTION

Technical analysis is by far the most popular approach to making short term directional forecasts in financial markets. However, technical analysis is not a well defined discipline, but rather an umbrella term for a very diverse collection of tools, ranging from formal statistical methods based on moving averages, to pattern recognition and magic numbers. The value of technical analysis remains controversial, with academic evaluations of technical trading rules divided over whether they can reliably yield excess risk-adjusted profits.

The aim of this chapter is to determine what tools are used by members of a small panel of technical analysts, how good their forecasts are, and whether there is any relation between forecast performance and the tools used.

Section 4.2 below describes briefly the main tools of technical analysis and reviews the ambivalent findings of recent academic research. Section 4.3 describes our survey and summarises main results concerning the use of different technical forecasting tools. Among other questions, the survey elicits information about the techniques used by the analysts, and 1-week ahead directional forecasts and likely trading ranges for three key foreign exchange rates and three stock indices. This allows us to test in Section 4.4 for the accuracy of the forecasts, their rationality, and the profitability of trading rules based on the forecasts. In Section 4.5 we look at the popularity of forecast techniques in the panel, and the association between forecasting techniques, accuracy and profitability. Some conclusions are drawn in Section 4.6.

4.2 TECHNICAL ANALYSIS

Technical analysis is “the study of market action, price, trading volume and open interest, primarily through the use of charts, for the purpose of forecasting future price trends” (Murphy, 1986). This is in contrast to Fundamental Analysis, which links movements in market prices to underlying economic factors such as interest rates, inflation and expected company profitability.

To illustrate the popularity of technical analysis, an internet search in January 2004 for keywords “Technical Analysis” turned up about 8.4 million sites, “Technical Analysis Stock Market” about 2.3 million sites, and “Technical Analysis Foreign Exchange” about 1.1 million. Corresponding searches for “Fundamental Analysis” produced less than half these numbers of sites (3.9 million, .7 million and .6 million respectively).

In an earlier survey of technical analysts, Taylor and Allen (1992) found that about 90% of their sample of around 150 foreign exchange dealers in London used technical analysis intensively for intra-day and 1-week ahead forecasts, and at least 75% use it in some way. Menkhoff (1997) confirmed this for a panel of fund managers in Germany, and Lui and Mole (1998) found the same picture in a survey of around 150 dealers in Hong Kong. Both surveys show that as the forecasting horizon increases, less weight is given to technical factors and more to fundamentals. There is also some evidence from the more recent surveys in the UK by Cheung, Chinn and Marsh (2000), in Japan and Singapore by Cheung and Wong (2000), and in the US by Cheung and Chinn (2001) that technical analysis is losing ground a little to fundamentals. However, for the majority of respondents technical analysis remains the dominant forecasting paradigm. Survey evidence on the use of technical analysis has focussed on the foreign exchange market, and there is no corresponding information covering stock markets, bond markets and commodity markets.

4.2.1 TECHNICAL ANALYSIS TOOLS


The tools of technical analysis are many and varied, and described in many get-rich-quick books aimed at professional traders and retail investors. A straightforward and detached description of the more popular techniques is given by Neely (1997), and Jobman (1998) gives detailed definitions and uses of a wide range of indicators. Technical indicators fall into two main categories. Some indicators are “statistical”, based on mathematical functions of recent prices. These indicators are easily replicable, and can be parameterised so as to maximise the probability of trading profitably. Other indicators are “judgmental”, and are based on the recognition of significant patterns in the sequence of prices. In some cases these patterns can be formalised. In others, the interpretation of the chart relies heavily on the experience and preferences of the analyst.

Figure 1 shows a typical chart, for the Yen/\$ exchange rate in 1998, overlaid with some of the most common statistical indicators. The chart shows daily open/ high/ low/close prices as vertical bars. A short term and long term backward moving average (MA) of closing prices has been drawn through the price series. If the closing price cuts the long term MA from above, or more conservatively if the short term MA cuts the long term MA from above, this would be interpreted as a SELL signal (and vice versa). In the Figure, we have arbitrarily used 5-day and 25-day moving averages. In practice, an analyst would choose windows which would have yielded consistent profits in the recent past, or reflected current market conditions.

The lower panel of Figure 1 plots another statistical indicator, the %K and %D “stochastics” for the Yen/\$ price series. These are examples of “oscillators”, designed to anticipate turning points in the market. Most technical analysis textbooks emphasise the importance of “confirmation”. A signal from one indicator should not be acted upon unless confirmed by others. So for example, a SELL decision might require both the moving average and oscillator to give a SELL signal. In the Figure, %K is the ratio of the difference between the current closing price and the low price

over the last 5 days, to the difference between the high and low of the past 5 days. %D is a 3-day moving average of %K. The rationale for this is that if the market is rising strongly, then each day is likely to close near a new high, so %K will be close to 1. As the market approaches a peak, the market may still rise, but fail to achieve new highs, so that %K will fall. A SELL signal occurs if %K (or more conservatively %D) has been high, but starts to fall through some critical “overbought” level (say 0.8). The length of window, and the overbought and oversold levels would again be chosen so as to optimise expected trading profits. Other oscillators used in the same way as %K are the Momentum Index, and the Relative Strength Index (RSI).

Figure 1. Technical Indicators – Statistical




Notes: The Stochastics %K is in black and %D is in red.

Figure 2 shows a selection of charts which might be used to support judgmental pattern recognition by technical analysts. The upper left panel shows the price series overlaid with “trendlines” – lines of support and resistance drawn through recent lows and highs respectively. If the price breaks below the support line, this might indicate a SELL. The rationale is that the market is moving to a lower trading range.

Conversely, if as in the Figure the price breaks above the resistance line, the signal is to BUY. Note that the lines need not be parallel. Often they are convergent, indicating increasing “price congestion”. They need not be a constant distance apart, nor need they be symmetric around the current price.

Figure 2. Technical Indicators – Patterns



The lower left panel shows the price series overlaid by Fibonacci retracement levels. In the Chart, the price has fallen from a significant peak to a trough, by about 20

Yen/\$. Some analysts believe that as the exchange rate recovers it will meet natural resistance levels at levels $20 \times F$ Yen above the trough, where F is a Fibonacci ratio (.236, .382, .50, .618, 1, 1.618, 2.618, 4.236, ..). Fibonacci numbers occur frequently in nature and in art and architecture, in the form of the “golden section”. There is no logical rationale for this belief about the financial markets, but it could be argued that recovery to Fibonacci-determined levels “looks right” on aesthetic grounds.

The upper right panel of Figure 2 shows the best known reversal formation, the “head-and-shoulders”. This occurs when a rising market suffers a reverse, recovers to a higher (peak) level, reverses again, and then in its next rally fails to rise higher than the previous peak and falls back. This is believed to herald a sustained fall in price, and so generates a SELL signal. Other reversal formations include “double tops” and “double bottoms”. In addition to these local turning point patterns, there is a long tradition of “wave theory”, including Dow Theory and the Elliott Wave theory. In these systems, the market is believed to rise and fall in a series of cycles, and its current position in the cycle determines the likelihood of a reversal, and natural support and resistance levels.

The lower right panel shows the price series as a “Candlestick Chart”. The daily open/close range is shown as a block, white if the price rose through the day, dark if it fell between the open and the close. This is a visually informative way of presenting the daily price data. Moreover some analysts believe that sequences of 1-, 2-, or 3-day patterns in these candlesticks can be used to trigger trading decisions. In the Figure, a long white candlestick occurs at the end of the upward price trend. This is followed by a black candlestick which opens higher than the previous day’s close, but closes down near the previous days open. This is a classic “dark cloud cover” pattern, which would indicate a SELL. This and many other exotically named candlestick patterns are described in Nison (1994).

Apart from these price-based indicators, some analysts – especially in equity and futures markets –also use indicators based on trading volume or hedging demand as revealed in figures for open interest.

4.2.2 EVIDENCE ON TECHNICAL ANALYSIS

Whether any of these indicators have value is hotly disputed. Their users evidently believe in their effectiveness, and there is a large market for the services of technical analysts, books on technical analysis, and information systems incorporating technical indicators. On the other hand, efficient markets theory implies that investors should not be able to persistently exploit patterns in data to generate excess risk-adjusted returns. Technical analysis is ridiculed by the respected academic Burton Malkiel in his classic commentary on the US stock market, “A Random Walk Down Wall Street” (Malkiel, 1997). For balance it should be noted that he treats fundamental analysis with only a little more respect.

In recent years some more measured studies have tried to assess the performance of technical analysis. Some of these have examined technical trading rules. Others have looked directly at the performance of technical analysts.

A number of early papers claim to test technical trading rules, but the rules they examine have never been popular among analysts. Among these are the classic studies of filter rules by Alexander (1961, 1964), Fama and Blume (1966), Sweeney (1986), and Levich and Thomas (1993). Filter rules involve buying if the price rises by more than some threshold percentage, and selling if it falls by more than this percentage. This constant and symmetric trading range is only loosely related to the concepts of resistance and support discussed above, and these studies should really be regarded as tests for a particular kind of market efficiency.

Brock, Lakonishok and LeBaron (1992) do test recognisable technical rules. Using daily closing prices of the Dow Jones Industrial Average Index in the years 1897-1986, they test the profitability of five moving average crossover rules, and breakouts from three trading ranges based on recent highs and lows. The moving average rules generally outperform naïve alternatives, including buy-and-hold. Subsequently, Sullivan, Timmermann and White (1998) test the best rules on a further 10 years data, and conduct a formal test for data-snooping biases. None of the rules performs well out of sample, suggesting a change in structure – possibly an increase in efficiency – in the US market.

Silber (1994) also looks at short/ long moving average crossover rules, but across a number of currency and commodity and interest rate futures markets. He finds the best rule for each market in one year, then uses this to trade in the following year. The rules prove profitable in the currency markets, but not in the commodity or interest rate futures markets. He ascribes this to price smoothing by central banks, which may reduce the ratio of noise to signal in the currency markets, hence reduce the risk that moving average traders are “whipsawed” – that is, forced into frequent trades in a non-trending market.

In two benchmarking papers, Osler (1998) and Osler and Chang (1999) conducted an exhaustive investigation of the “Head and Shoulders” pattern in 100 shares traded on the NYSE in the years 1962-93, and in six currency markets in the years 1973-1994. With respect to the currencies, three showed a profit and three showed a loss, and none would have outperformed naïve alternative forecasting rules. Similarly, there was no evidence that the Head and Shoulders top presaged a fall in share prices. In many cases it was indeed followed by a fall, but in as many other cases it was followed by a rising price.

Each of these studies focusses on a single type of trading rule. In practice, analysts generally seek “confirmation” of trading signals by looking at more than one indicator. Some studies examine whether this would be profitable, though again the

results mixed. Pruitt and White (1988), for example, report positive results from combining moving average, relative strength and trading volume indicators for stock market timing. On the other hand, Neely, Weller and Dittmar (1997) have difficulty in finding combinations of technical rules which produce positive risk-adjusted returns to currency trading, in spite of using an advanced “genetic algorithm” search procedure.

The other problem with the above studies is that they cannot credibly claim to mimic the behaviour of technical analysts. Analysts may give weights to the different rules which vary in complex ways over time, and indeed may well combine technical and fundamental information.

The earliest study to look directly at the performance of a technical analyst rather than a technical rule is Cowles (1933, 1944) examination of the track record in the years 1902-29 of W.P. Hamilton, the author of “Dow Theory”. This data has recently been revisited by Brown, Goetzmann and Kumar (1998). Though Cowles was initially sceptical, the balance of evidence now suggests that Hamilton’s ability to call stock market rises and (especially) falls was significantly better than chance.

Murphy (1986) looked not at the recommendations but at the trading records of 11 technical futures funds tracked by monthly magazine “Commodities” through the years 1980-85. Most funds earned positive returns and all outperformed a money market benchmark and the S&P500 index. However, on a risk-adjusted basis they did not fare so well, only just outperforming the index. Whether this amounts to a criticism of technical analysis is a moot point, since most actively managed funds tend over the long run to underperform passive index tracking portfolios.

In conjunction with their survey of the use of technical analysis in the London foreign exchange markets, Allen and Taylor (1990) and Taylor and Allen (1992) track one week and one month ahead forecasts of about 20 analysts. Only one outperformed the no-change “Random Walk” forecast, though most outperformed more complicated

time series models. Each week there was a wide range of disagreement among analysts about the likely direction of change and target rate for the currencies. This lack of consensus suggests that technical factors are unlikely to be “self-fulfilling”, and that they are unlikely to be destabilising in the sense of causing runs up or down in exchange rates.

Curcio and Goodhart (1992) and Osler (2000) test the value of published data on support and resistance levels in major currency markets. In the first study, the data are from the Reuters information service in the period April-June 1989. The authors test the profitability of buying/ selling when a resistance/ support line is broken, with positive results. In the second study, a much larger data set consisting of daily data from six banks and information services in the period January 1996-March 1998 is used. The author tests whether rates tend to cluster around, and bounce back from the published levels. Again, the results are positive. Moreover, there is a reasonable degree of agreement among institutions about where the critical levels are. This suggests that – contrary to the conjecture of Taylor and Allen (1992) - the use of trendlines by technical analysts may impose some self-fulfilling patterns on exchange rates, in the form of persistence of prices around support and resistance levels.

4.3 OUR SURVEY

The panel for our survey consisted of 20 technical analysts. Of these, 15 were employed by the same leading international financial information and analysis service, Standard & Poor's MMS. The rest worked at leading international investment banks. The analysts were assigned letter codes, and their affiliations, job titles and locations are shown in Table 1. All respondents are experienced, having worked as technical analysts for between 2 and over 10 years. Most are involved in providing daily and weekly on-line technical commentary across wire services such as Reuters and Telerate, to dealers in currency markets, bond markets and stock markets. Most

are located in London, but a few are based in New York, San Francisco and Singapore.

Table 1. Survey panel

Code	Company	Job Title
A	Standard & Poor's MMS Singapore	Senior Technical Analyst
B	Standard & Poor's MMS Singapore	Technical Analyst
C	Standard & Poor's MMS Singapore	Technical Analyst
D	Standard & Poor's MMS Singapore	Technical Analyst
E	Overseas Union Bank Singapore	Assistant manager
F	Standard & Poor's MMS London	Technical Analyst
G	Standard & Poor's MMS London	Manager Technical Analysis
H	Standard & Poor's MMS London	Director Technical Analysis
I	Standard & Poor's MMS London	Head of Technical Research
J	Standard & Poor's MMS London	Manager Technical Analysis
K	Standard & Poor's MMS London	Technical Analyst
L	Standard & Poor's MMS London	Senior Technical Analyst
M	Standard & Poor's MMS London	Senior Technical Analyst
N	Bank of American Express London	Senior Director
O	Merrill Lynch London	Senior Technical Analyst
P	Deutsche Morgan Grenfell London	Senior Technical Analyst
Q	Royal Bank of Scotland London	Senior Technical Analyst
R	Standard & Poor's MMS US	Technical/Options Strategist
S	Standard & Poor's MMS US	Senior Technical Analyst
T	Standard & Poor's MMS US	Senior Technical Analyst

The immediate objective of the survey is to determine how accurate the predictions of this set of analysts are, what methods they use, and whether there are associations between forecast quality and methods used. Our sample is not a random or representative one, so we are limited in how far results on the panel can be used to generalise about the whole population of technical analysts. However, since our analysts are experienced professionals receiving regular daily feedback from their forecasts, we would expect their performance to be above-average for the analyst population.

The survey started on 9 January 1998, and was terminated on 6 November 1998, giving data for 44 weeks responses. Every Friday afternoon around the London market close of 15:30 GMT we faxed a questionnaire to all 20 participants in the survey. Participants were requested to return the survey before 11:00 GMT on Monday. Most respondents respected these deadlines, some replying as early as Friday evening. One survey was sent on Thursday on 9th of April 1998 due to the Easter Bank holiday on Friday. Similarly, five were sent on Bank Holiday Mondays, with a deadline of 11:00 GMT on Tuesday.

Holidays and absence meant that even the most assiduous respondents did not complete questionnaires for all 44 weeks. As might be expected the highest number of responses were achieved early in the life of the survey, with an average of 14 participants in the early (January – March) surveys, but only 9 in the later (September–November) surveys. The maximum number was 17, in the week ending 23 January. The minimum was 5 in the week ending 2 October. Of the initial 20 respondents, 14 analysts stayed in the survey for more than 20 weeks. Our analysis of survey responses focuses on these 14 individuals. Their individual response rates are detailed in Table 2.

Table 2. Survey response frequency

Code	A	B	F	G	H	I	J	K	L	M	N	P	S	T
DEM	36	32	41	36	35	40	42	37	39	20	21	26	24	25
JPY	36	32	42	36	35	41	42	37	39	20	22	27	24	26
GBP	36	31	42	36	35	41	42	37	39	20	22	27	24	25
DJIA	36	32	41	36	35	40	42	37	39	20	21	26	24	25
FTSE	0	0	28	36	35	33	42	37	39	20	21	26	11	12
NIKKEI	36	31	42	36	35	40	41	37	39	20	21	25	24	26

Notes: Table shows number of weeks responses by currency/ stock index for the 14 most regular respondents out of maximum of 44 weeks.

The first page of the questionnaire is reproduced as Table.3. This page asks about likely trends in the Deutschemark/US Dollar exchange rate. Five other pages asked identical questions about the Yen/ Dollar and Dollar/ Sterling exchange rates, and about the Dow Jones, Nikkei and FTSE100 stock market indexes. The survey provides a base level price for each exchange rate and index, in the form of the latest London closing price. Some analysts specialise in either currency markets or equity markets, or in one particular region. So even in weeks when they did reply, the 14 regular respondents might not complete the questionnaire for all six target variables. For example, forecasters A and B did not provide any forecasts of the FTSE100 index.

Table 3. Sample survey questionnaire

USD-DM = 1.7835 16:15 GMT London			Horizon:	
			1-week	1-month
1. What is your target rate 1-week/ 1 month ahead?				
2. What are your target ranges?			High	
			Low	
3. What is the likely trading tone? tick <u>ONE</u>				
Trending				
Consolidating				
Correcting				
Reversal				
4. Please give two crucial support/ resistance levels:			1st	2nd
Support:				
Resistance:				
5. Which technique s gave you the main signal for your 1-week forecast? check <u>UP TO 3</u> techniques				
Moving Averages	Trendlines	Elliot Wave	Market Profile	
Momentum	Channels	Gann Analysis	Congestion Breakout	
Stochastics	Cycles	Fibonacci Numbers	Volume/ Open Int.	
Relative Strength	Seasonal Pattern	Candlestick Charts	Other specify :	

Question 1 asks for 1-week and 1-month ahead point forecasts of the target variable, and Question 2 for quantitative estimates of the likely 1-week and 1-month market highs and lows. All the analysts who did respond provided these quantitative forecasts, though as we have seen technical analysts are really more concerned with directional or qualitative predictions. Question 3 asks whether the market is likely to continue its current trend (trending), change direction (reversal), move sideways (consolidating) or return to a previous trend (correcting). This may also be significant for the choice of forecasting technique, since some methods such as moving averages are useful only in a trending market. Technical analysts are also concerned with trader support, and Question 4 asks for critical resistance and support levels at which traders might be expected to place conditional limit and stop-loss orders. Finally, question 5 asks what techniques the analyst used to make his or her 1-week forecast. The language of the questionnaire and the response categories were selected after discussion with the analysts, and so differ a little from the terminology used in our survey of academic work above.

4.4 SURVEY RESULTS: ACCURACY, RATIONALITY AND PROFITABILITY

We start by looking at the accuracy of individual and group forecasts, then examine their rationality and calibration, and the profitability of trading on the directional forecasts.

4.4.1 FORECAST ACCURACY

Tables 4-6 compare the track records of the 14 regular forecasters, using three common error metrics – directional accuracy (DA), mean absolute percentage error (MAPE) and root mean square percentage error (RMSPE). In addition to the

individual analyst figures we have also calculated the accuracy of the “Consensus Forecast”, defined as the average of the individual point forecasts from the survey.

Directional accuracy is defined as the fraction of weeks in which the analysts correctly predicted the direction of change in the exchange rate or stock index. The DA metric is particularly relevant since the objective of technical analysis is to provide forecasts of turning points, to inform buy and sell decisions. Note that a DA of 0.5 would be expected from a forecaster with no skill at forecasting the direction of change. The figures show that, as expected, the one week forecasts were more accurate than the one month forecasts. For the currencies only 1 out of 42 1-week forecasts showed a directional accuracy below 0.50, whereas 14 out of 42 1-month forecasts were below this no-skill benchmark. The figures also show that some target variables were easier to predict than others. The Deutschemark, for example, was harder to forecast than Sterling, and both were harder than the Yen. The stock indexes were rather harder to predict than the currencies. The one-week ahead Consensus forecasts for sterling, for example, were correct in 73% of weeks, but the corresponding forecasts for the FTSE100 index were correct in only 50% of weeks. Overall, the impression left by Table 4 is that the analysts were relatively successful at predicting currency movements. In 26 out of 42 cases they had a directional accuracy in excess of 0.6. They were less successful with the stock indexes, but even here a majority of forecasters achieved a DA above 0.5. There are of course two possible reasons why the equity forecasts may be less accurate than the currencies. One is that the participants from Standard & Poor's MMS tend to be more specialised in currency and fixed income markets, rather than in equities. Another is that the equity markets may simply be more difficult to forecast.

Table 4. Directional Accuracy

1-week horizon

Forecaster	Max Weeks	Currencies			Stock Indexes		
		DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	0.56	0.67	0.67	0.39	0.58	
B	32	0.66	0.72	0.65	0.50	0.52	
F	42	0.62	0.71	0.62	0.56	0.62	0.43
G	36	0.53	0.56	0.64	0.64	0.61	0.44
H	35	0.54	0.66	0.54	0.46	0.40	0.60
I	41	0.61	0.66	0.68	0.68	0.60	0.52
J	42	0.48	0.62	0.64	0.64	0.56	0.40
K	37	0.51	0.59	0.62	0.57	0.57	0.51
L	39	0.51	0.67	0.64	0.62	0.51	0.44
M	20	0.60	0.55	0.60	0.45	0.50	0.50
N	22	0.50	0.59	0.59	0.38	0.52	0.33
P	27	0.59	0.70	0.85	0.62	0.56	0.42
S	24	0.58	0.79	0.54	0.54	0.50	0.45
T	26	0.54	0.81	0.50	0.56	0.62	0.50
Consensus	44	0.55	0.70	0.73	0.55	0.57	0.50

1-month horizon

Forecaster	Max Weeks	Currencies			Stock Indexes		
		DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	0.64	0.69	0.44	0.50	0.53	
B	32	0.44	0.75	0.71	0.69	0.58	
F	42	0.50	0.57	0.45	0.68	0.57	0.57
G	36	0.64	0.50	0.56	0.36	0.50	0.33
H	35	0.51	0.54	0.37	0.58	0.60	0.44
I	41	0.49	0.46	0.59	0.50	0.50	0.39
J	42	0.45	0.69	0.50	0.62	0.41	0.50
K	37	0.41	0.59	0.49	0.59	0.54	0.49
L	39	0.59	0.49	0.54	0.59	0.46	0.42
M	20	0.45	0.50	0.40	0.70	0.50	0.50
N	22	0.59	0.59	0.77	0.38	0.43	0.48
P	27	0.48	0.78	0.33	0.54	0.68	0.54
S	24	0.50	0.83	0.50	0.54	0.38	0.55
T	26	0.62	0.54	0.54	0.48	0.46	0.33
Consensus	44	0.59	0.64	0.55	0.5	0.52	0.43

Table shows proportion of weeks in which forecaster correctly predicted the direction of change of the target variable over 1 week and 1 month horizon. Figures in bold are significantly greater than 0.5 at the 95% significance level.

Table 5 shows corresponding figures for the mean absolute percentage error (MAPE) in the analyst forecasts, and in the Consensus forecast, and Table 6 shows figures for the related root mean square percentage error (RMSPE) measure. The RMSPE penalises large errors disproportionately, so produces some differences in ranking. However, the general pattern of both MAPE and RMSPE figures are the same, so we discuss both sets of data together.

As a benchmark in both tables we show the MAPE for a naïve "Random Walk" forecast, based on the Friday closing price at each survey date. Previous evidence has shown that currency and interest rate forecasters are generally dominated by the no-change, Random Walk forecast forecasts (Boothe, 1983; Boothe and Glassman, 1987; MacDonald and Hein, 1989; Kolb and Stekler, 1996). This is not so true of our panel of technical analysts. At the 1-week horizon, for example, the Consensus forecast is more accurate than the Random Walk for all currencies and stock indexes except for the FTSE100. The consensus forecast is inevitably more accurate than most of the individual forecasts, however, and for most series only between 1 and 6 individuals outperform the Random Walk. An exception is the Yen, where the large fall early in September 1998 was forecast by most analysts, giving them an advantage over the naïve forecast. Performance deteriorates in absolute terms and relative to the benchmark, as we move from the 1 week to 1 month horizon. However, the Nikkei index is the only series where the Random Walk model is unambiguously better even at the one-month horizon.

Table 5. Mean absolute percentage errors

MAPE 1-week horizon

Forecaster	Max weeks	Currencies			Stock Indexes		
		DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	1.24	1.82	0.93	2.74	3.23	
B	32	1.00	1.67	0.96	2.32	3.33	
F	42	1.10	1.82	0.93	2.16	3.03	2.37
G	36	1.29	2.40	1.01	2.64	3.80	3.63
H	35	1.22	2.23	1.19	2.43	3.64	2.49
I	41	0.96	2.10	1.00	2.72	3.29	2.82
J	42	1.11	1.82	0.86	2.43	3.42	3.18
K	37	1.30	1.95	1.06	2.96	4.04	2.84
L	39	1.13	1.47	0.99	1.99	3.22	2.82
M	20	1.13	1.61	0.88	2.03	3.23	2.14
N	22	1.03	2.10	0.99	2.58	5.08	4.06
P	27	1.24	1.99	0.99	2.61	4.19	4.20
S	24	0.95	1.28	0.99	2.32	2.96	2.45
T	26	1.15	1.29	1.02	2.24	2.61	1.80
Consensus	44	0.89	1.67	0.83	1.96	2.59	2.51
Random Walk	44	1.04	2.23	0.87	2.05	2.85	2.14

MAPE : 1-month horizon

Forecaster	Max weeks	Currencies			Stock Indexes		
		DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	1.87	3.63	1.75	6.00	7.64	
B	32	2.58	3.66	1.43	3.87	7.06	
F	42	1.98	5.01	1.77	4.91	5.98	3.80
G	36	2.66	6.86	2.79	8.22	10.43	8.40
H	35	2.41	4.03	2.06	5.23	6.08	5.43
I	41	1.92	4.99	1.79	6.63	7.12	5.52
J	42	2.93	4.85	2.18	15.30	14.04	14.04
K	37	3.00	4.56	1.59	6.56	7.61	7.17
L	39	2.14	5.26	1.88	5.58	6.59	6.38
M	20	2.02	3.48	2.02	3.93	5.85	3.80
N	22	1.47	3.38	1.38	6.76	9.25	9.70
P	27	2.74	4.22	2.14	5.65	7.62	6.87
S	24	2.39	3.98	1.71	5.42	6.66	3.56
T	26	3.43	4.39	2.21	5.13	5.69	4.78
Consensus	44	2.05	4.16	1.45	5.61	5.88	6.12
Random Walk	44	2.01	4.25	1.46	5.03	5.25	5.08

Notes: Table shows mean absolute percentage error in 1 week and 1 month ahead forecasts.

Table 6. Root mean square percentage errors

RMSPE : 1-week horizon

Forecaster	Max weeks	Currencies			Stock Indexes		
		WGM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	1.70	2.35	1.17	3.55	3.74	
B	32	1.39	2.34	1.20	3.02	4.40	
F	42	1.39	2.86	1.20	2.70	3.97	2.97
G	36	1.60	3.38	1.31	3.45	4.95	4.54
H	35	1.67	2.91	1.47	2.98	4.58	3.13
I	41	1.19	2.98	1.32	5.71	4.11	3.50
J	42	1.40	2.42	1.13	3.03	4.44	3.80
K	37	1.56	2.41	1.40	3.65	4.90	3.59
L	39	1.47	1.76	1.25	2.58	4.14	3.62
M	20	1.64	2.01	1.23	2.48	4.35	2.77
N	22	1.36	2.50	1.18	3.07	6.24	4.82
P	27	1.74	3.53	1.15	3.73	5.17	5.02
S	24	1.19	1.63	1.25	2.90	3.98	3.20
T	26	1.51	1.60	1.21	3.03	3.37	2.43
Consensus	44	1.19	2.58	1.06	2.53	3.42	3.23
Random Walk	44	1.26	3.37	1.10	2.85	3.68	2.78

RMSPE : 1-month horizon

Forecaster	Max weeks	Currencies			Stock Indexes		
		DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	36	2.66	5.72	2.16	8.12	9.89	
B	32	3.39	5.58	1.76	5.11	8.51	
F	42	2.50	6.19	2.20	6.43	7.69	4.84
G	36	3.13	8.37	3.40	9.99	13.04	9.96
H	35	3.21	5.21	2.44	6.72	7.70	6.46
I	41	2.35	5.72	2.14	9.07	8.74	7.28
J	42	3.56	6.64	2.59	18.88	16.17	20.10
K	37	3.69	6.08	2.04	8.25	9.10	9.05
L	39	2.53	6.81	2.26	6.81	7.94	7.80
M	20	2.59	4.33	2.54	5.04	7.83	4.86
N	22	1.96	4.21	1.87	7.62	11.76	11.19
P	27	3.50	5.12	2.56	7.85	9.72	8.68
S	24	3.26	6.04	2.24	6.81	7.70	4.28
T	26	4.58	6.12	2.72	6.56	7.16	5.46
Consensus	44	2.60	5.66	1.79	6.82	7.65	7.57
Random Walk	44	2.59	5.76	1.74	6.21	6.67	6.33

Notes: Table shows root mean square percentage error in 1 week and 1 month ahead forecasts.

Tables 4-6 reveal differences among forecasters in accuracy as averaged over the weeks of the survey. It is natural to ask whether these differences are statistically significant, given variations in performance from week to week. This problem has been addressed by Batchelor (1990) for the case where the forecast panel is balanced, in the sense that all n respondents make the same number of predictions, for each of the T weeks of the survey. Given the sum of ranks r_i for forecaster i , to test equality in forecaster ranks he calculates the Friedman (1937) analysis of variance by ranks statistic

$$g = \sum_{i=1}^n \frac{\{r_i - T \cdot (n+1)/2\}^2}{Tn(n+1)/12}$$

which is $\chi^2(n-1)$ under the null of no difference in average rank. Here we generalise the method to the case where not all forecasters make forecasts each week. Let the number of forecasts made in week t be n_t so that the average rank in week t is $(n_t+1)/2$, and let $d_{it} = 1$ if forecaster i makes a forecast in week t . Then the test statistic becomes

$$h = \sum_{i=1}^n \frac{\{r_i - \sum_t d_{it} (n_t + 1)/2\}^2}{\sum_t d_{it} n_t (n_t + 1)/12}$$

which is again $\chi^2(n-1)$ under the null of no difference in average rank. The results are summarised in Table 7, using MAPE (= RMSPE) ranks. Only in the case of the Dow and the FTSE can we conclude that the differences in forecast accuracy across individual recorded in Tables 5 and 6 are statistically significant. There are no significant differences across accuracy rankings for the three currencies or for the Nikkei index.

Table 7. Tests for equality of accuracy ranks across forecasters

	<i>h</i> -statistic	<i>p</i> -value
DEM	9.40	0.74
JPY	11.84	0.54
GBP	13.48	0.41
DJIA	26.72	0.01
NIKKEI	14.51	0.34
FTSE	28.17	0.00

Notes: Table shows Batchelor (1990) test statistics and associated p-values for test of equality of MAPE (RMSPE) ranks across our panel of forecasters.

4.4.2 FORECAST RATIONALITY

A forecast is defined as rational in the sense of Muth (1960) if it utilises relevant information in an optimal way. Many empirical tests have been conducted to establish the rationality of survey-based forecasts. Many of these have focussed on macroeconomic aggregates like inflation and GDP growth, and occasionally on interest rate forecasts (Batchelor and Dua,1992; Chinn et al,1991; Keane and Runkle,1990).

The basic form of rationality test is a regression of the form

$$A_{t+1} - F_{it} = a_i + b_i X_{it} + u_{it} \quad (1)$$

where A_{t+1} is the realised value of the target variable following the forecast made in week t , F_{it} is the forecast made by forecaster i in week t , and X_{it} is some information known to forecaster i at the time the forecast is made. Rationality would be rejected if the coefficient b_i was significantly non-zero. The idea is that if the relationship (1) is stable and knowable to the forecaster at t (both contentious issues in practice), then the forecast error could have been reduced by making a linear adjustment to the

forecast based on the value of X_{it} . Since forecasters are certain to know the current value of the target variable A_t when making their forecast – recall that this is given to them on the survey questionnaire – a minimal condition for rationality is that $(a_i, b_i) = (0,0)$ in

$$A_{t+1} - F_{it} = a_i + b_i A_t + u_{it} \quad (2)$$

This test equation is more usually estimated as the regression between the actual change in the target variable and the forecast change:

$$A_{t+1} - A_t = \alpha_i + \beta_i (F_{it} - A_t) + u_{it} \quad (3)$$

where $\beta_i = 1 - b_i$ so that the so-called “unbiasedness condition” for rationality becomes $(\alpha_i, \beta_i) = (0,1)$.

Table 8 shows the results of this test for 1 week ahead forecasts, by forecaster and target variable. The consensus forecast is noticeably more rational than the individual forecasts, with unbiasedness only rejected for the FTSE. Most forecasters produce rational forecasts only for the Yen, and only forecaster I manages to produce 3/6 rational forecasts.

Table 8. Rationality of 1-week ahead forecasts

	DEM	JPY	GBP	DJIA	NIKKEI	FTSE
A	30.76	4.44	14.95	23.74	25.36	
B	9.55	1.85	14.56	9.76	16.11	
F	14.94	0.47	10.96	9.16	17.63	27.13
G	24.94	4.24	25.06	24.71	33.89	67.81
H	2.29	5.53	25.63	5.13	16.93	9.48
I	3.15	0.11	25.18	3.90	19.54	27.22
J	11.30	9.43	6.63	15.50	19.93	36.46
K	14.27	9.08	20.26	21.46	38.44	23.77
L	13.11	1.85	18.44	6.74	12.27	20.83
M	16.63	3.40	5.26	43.99	8.76	25.83
N	6.02	3.36	7.58	21.98	34.34	44.39
P	18.27	2.55	24.58	25.83	37.60	39.83
S	11.78	0.29	12.54	3.79	16.57	
T	20.82	3.23	7.52	6.24	6.57	
Consensus	3.13	4.84	6.00	1.71	4.20	16.27

Notes: Table shows test statistics for the “unbiasedness test” for rationality by forecaster, $\alpha_i, \beta_i = 0, 1$ in the regression $A_{t+1} - A_t = \alpha_i + \beta_i F_{it} - A_t + u_{it}$ where A = actual value of target variable, F = forecast value. The statistic is distributed as χ^2_2 under the null of rationality, with 5% critical value 5.99. Cases where unbiasedness is not rejected as shown in **bold**.

4.4.3 PROFITABILITY

The aim of the analysts in our survey is to give profitable trading signals to their clients. While mean-square accuracy and forecast rationality is of some academic interest, the criterion by which the analysts are judged is whether they can help their clients make money. Studies by Boothe and Glassman (1987) and Leitch and Tanner (1991) suggest that, in exchange rate and interest rate forecasting there is little correlation between profitability and accuracy. The reason is that changes in prices in financial markets typically experience extreme rises and falls more often than would be expected under a normal distribution. Getting on the right side of the market ahead of these “big hits” is more important for profitability than minimising error variance, or even maximising directional accuracy. The sets of forecasts used in these studies

are very limited. For example, Leitch and Tanner use only two analyst-based forecasts, the forward rate, and three artificially generated ex post time series predictions. Our data let us test their results using a larger panel of real-time forecasts.

To determine the profitability of a set of forecasts, some trading rule must be assumed. Here we make the simplest assumption, that the trader takes a long position when the forecast is for the target to rise, and takes a short position if the forecast is a fall. This presumes that stock market investors can take short positions. The size of the position taken each week is assumed to be the same regardless of the size of the forecast change in market price.

Table 9A summarises the profits to a dollar-based investor made in currency markets by following the individual analysts and the consensus forecast. We assume that trading is on a rolling spot basis so that investors are credited with any weekly interest differential in addition to gains on spot rates. Short term interest rates in the US were around 5.5% in 1998, and about 0.5%, 3.5% and 7.5% in Japan, Germany and the UK. As benchmarks we show a “buy and hold \$” profit, which would be achieved from a series of 1-week US money market deposits covering all weeks of the survey. We also show a “buy and hold FX” benchmark, representing the dollar returns on a foreign currency deposit held through the survey period. The table also shows the standard deviation of weekly percentage returns, and the Sharpe ratio – the ratio of the excess return of each analyst over the benchmark buy and hold \$ return, normalised by the standard deviation. The Sharpe ratio is a common measure of risk-adjusted investment performance.

Table 9A. Profitability of trading on 1-week targets: currencies

Forecaster	DEM				JPY				GBP			
	av profit %/ week	std dev %/ week	Sharpe ratio	Total Profit	av profit bps	std dev bps	Sharpe ratio	Total Profit	av profit bps	std dev bps	Sharpe ratio	Total Profit
A	0.00	1.24	-0.09	99.7	1.40	2.25	0.58	163.8	0.43	1.01	0.32	116.6
B	0.32	1.25	0.17	110.6	1.01	2.37	0.38	136.8	0.35	1.06	0.23	111.3
F	0.29	1.17	0.16	112.6	1.30	2.86	0.42	169.2	0.34	1.12	0.21	115.2
G	0.43	1.27	0.26	116.5	1.15	3.17	0.33	148.7	0.17	1.15	0.06	106.2
H	0.29	1.32	0.14	110.4	0.95	2.66	0.32	137.8	0.18	1.17	0.06	106.1
I	0.48	1.23	0.30	121.2	1.28	3.08	0.38	165.6	0.48	1.08	0.35	121.4
J	-0.01	1.27	-0.09	99.1	0.90	3.19	0.25	143.0	0.46	1.09	0.32	120.9
K	0.24	1.33	0.10	109.1	0.72	2.40	0.26	129.0	0.42	1.17	0.27	116.3
L	0.22	1.31	0.08	108.4	1.21	2.22	0.50	158.6	0.48	1.03	0.36	120.4
M	0.28	1.21	0.14	105.5	0.06	1.97	-0.02	100.9	0.66	1.08	0.52	114.0
N	0.21	1.25	0.09	104.6	0.60	2.32	0.21	113.3	0.45	1.15	0.30	110.3
P	0.31	1.38	0.15	108.5	0.70	3.58	0.17	118.7	0.54	1.04	0.41	115.4
S	0.41	1.20	0.26	110.2	1.15	2.06	0.51	131.1	0.32	1.02	0.21	107.9
T	0.19	1.16	0.08	105.0	1.03	1.95	0.47	129.9	0.05	1.12	-0.05	101.0
Consensus	0.33	1.24	0.18	115.2	1.49	2.89	0.48	188.6	0.62	1.01	0.51	131.1
Buy/Hold \$	0.11	0.00		104.8	0.11	0.00		104.8	0.11	0.00		104.8
Buy/Hold FX	-0.06	1.26	-0.13	97.0	0.04	3.23	-0.02	99.5	0.25	1.12	0.13	111.5

Notes: Table shows for each currency the average weekly percentage profits in \$ and the standard deviation of weekly percentage profits on trades based on analyst directional forecasts in the 44 weeks of the survey. Profits include interest rates. Benchmarks are Buy/Hold \$ = profits from a rolling 44 week US money market deposit, and Buy/Hold FX = dollar profits from a 44 week rolling foreign currency investment. The Sharpe ratio is analyst profit-buy/hold \$ / standard deviation of analyst profit, a measure of risk-adjusted returns to a dollar-based investor.

Almost all analyst forecasts give rise to profits in excess of the buy and hold benchmarks. The scale of profits is very striking for the Yen, and slightly higher for Sterling than for the Deutschemark. The Consensus forecast performs consistently well, and trading based on the consensus yields higher profits than most individual forecasters. Among individual analysts there are obvious differences currency by currency – I for example is much better than J in predicting the Deutschemark – but relative performance is not consistent across currencies, and rank correlations are if anything slightly negative. So a superior performance in one market is no guarantee of a superior performance in another.

Table 9B shows corresponding profitability statistics for forecasts of stock indices. Here we assume that analysts can take long or short positions in the index, and so receive or pay away dividends. The dividend yield in the US averaged 1.5% in 1998, and was around 1% in Japan and 2.5% in the UK. The profits in this case are calculated from the point of view of a local currency based investor. The benchmarks are therefore a series of 1-week money market deposits, and alternatively a continuous long position in the index.

Performance in stock market trading is markedly inferior to that for currency trading. Only four analysts (a different four in each case) beat the buy-and-hold-index benchmark. Though all but 1 analyst beat the Nikkei index, only 4 beat the Dow, and only 2 analysts beat the FTSE. Even the consensus forecast failed to beat the FTSE, and even for the other markets it is less dominant than in the case of currency markets. Correlations of profitability across markets is low, suggesting that expertise in one market does not transfer to others. An important exception is the rank correlation of 0.5, between profits on the Yen and profits on the Nikkei, showing that forecasters A and B (both in Singapore), and F and I, may have some specialist insight into the Japanese markets.

Table 9B. Profitability of trading on 1-week targets: stock indices

Forecaster	DJIA				NIKKEI				FTSE			
	av profit %/ week	std dev %/ week	Sharpe ratio	Total Profit	av profit %/ week	std dev %/ week	Sharpe ratio	Total Profit	av profit %/ week	std dev %/ week	Sharpe ratio	Total Profit
A	-0.14	2.83	-0.09	93.7	1.30	3.45	0.37	156.1				
B	0.54	2.82	0.15	117.5	0.73	3.49	0.21	123.2				
F	0.69	2.76	0.21	130.5	0.87	3.56	0.24	140.0	-0.17	2.45	-0.13	94.6
G	0.62	2.66	0.19	123.6	0.67	3.55	0.18	124.3	-0.69	2.60	-0.32	77.1
H	0.61	2.95	0.17	122.0	-0.51	3.76	-0.14	81.7	0.70	3.03	0.18	125.5
I	1.25	2.68	0.43	162.1	0.78	3.41	0.23	133.3	-0.11	2.71	-0.09	95.4
J	0.96	2.83	0.30	147.0	0.57	3.57	0.16	123.3	-0.25	2.79	-0.14	88.4
K	0.69	2.98	0.20	127.1	0.47	3.44	0.13	116.6	0.77	2.93	0.21	131.0
L	1.16	2.78	0.38	154.4	0.50	3.74	0.13	118.3	-0.02	2.97	-0.06	97.4
M	-0.07	1.56	-0.11	98.3	-0.16	3.64	-0.05	95.7	-0.11	1.96	-0.13	97.6
N	-0.05	2.07	-0.08	98.4	0.08	4.01	0.02	100.0	-0.66	2.85	-0.28	86.3
P	1.67	2.89	0.54	152.2	1.40	4.00	0.35	139.1	-0.17	3.21	-0.10	94.4
S	0.32	2.95	0.07	107.0	0.49	3.37	0.14	111.0				
T	0.60	2.96	0.17	114.9	0.91	3.01	0.30	125.1				
Consensus	0.91	2.80	0.29	146.5	0.58	3.58	0.16	125.3	-0.06	2.84	-0.07	95.7
Buy/Hold Cash	0.11	0.00		104.8	0.01	0.00		100.4	0.14	0.00		106.5
Buy/Hold Index	0.67	2.86	0.20	131.7	-0.10	3.62	-0.03	93.0	0.20	2.85	0.02	107.1

Notes: Table shows for each market the average weekly percentage profits in local currency and the standard deviation of weekly percentage profits on trades based on analyst directional forecasts in the 44 weeks of the survey. Profits include dividend yields, and investors are assumed to be able to take a short position in the market. Benchmarks are Buy/Hold Cash = profits from a rolling 44 week money market deposit, and Buy/Hold Index = profits from a 44 week rolling foreign currency investment. The Sharpe ratio is analyst profit-buy/hold cash / standard deviation of analyst profit, a measure of risk-adjusted returns.

To illustrate the value of the analyst currency forecasts, Figure 3A shows notional trading profits in dollars from switching into Deutschmarks when the Consensus forecast is for the dollar to depreciate and vice versa. The figure shows clearly that the superior performance of the Consensus forecast would have yielded profits well in excess of a benchmark strategy of buying and holding dollars. To give some idea of the range across individuals, we also show profits from trading on the basis of the forecasts of Forecaster I (one of the most accurate) and Forecaster J (less accurate than the Consensus). These two individuals were used because they made forecasts in almost every week of the survey. The more accurate forecaster would also have earned higher profits than the Consensus, and the less accurate forecaster would have earned less. Interestingly, even the less accurate forecaster produces profits in excess of the simple buy and hold strategy.

Figure 3B shows the track of profits from trading the Dow, on the basis of the consensus forecast and forecasters I and A (below average – see Table 9B). In this case the Consensus and analyst I do end the period higher than the Dow, but for half of the year they underperform, and really only overtake the index by correctly calling its decline in July and August 1998.

Figure 3A. Profitability of trading on technical forecasts: Deutschemark-Dollar





Figure 3B. Profitability of trading on technical forecasts: DJIA



The loose connection between commercial value and accuracy is underlined by Table 10, where we report rank correlations (across forecasters) between profitability and directional accuracy, mean absolute percentage forecast error, and root mean square percentage error (the latter two ranked from lowest to highest). There are significant positive correlations between profitability and directional accuracy, especially for the stock indices. However, correlations between profitability and MAPE and RMSPE are very low, and negative in the case of the Dow and FTSE. Our data therefore confirm for a larger sample of forecasters the findings of Boothe and Glassman (1987) and Leitch and Tanner (1991) that conventional error metrics are not reliable guides to the profitability of trading on analysts forecasts.

Table 10. Rank correlations between accuracy and profitability

Correlation With	Currencies			Stock Indexes		
	DEM	JPY	GBP	DJIA	NIKKEI	FTSE
Directional Accuracy	0.59	0.41	0.52	0.83	0.75	0.73
MAPE	0.27	0.10	0.55	-0.11	0.17	-0.43
RMSPE	0.30	0.01	0.33	-0.31	0.35	-0.48

Notes: Rank correlations across 14 forecasters 12 for FTSE between profitability Table 8A & 8B and direction accuracy Table 4, mean absolute percentage forecast error Table 5, and root mean square percentage error Table 6.

4.5 USE OF TECHNICAL INDICATORS

We now move to the question of how the analysts made their forecasts. Table 11 summarises the number of times each technique was mentioned in connection with each market. The techniques are ranked in the table according to popularity, as measured by the total number of weeks when the technique was mentioned. The most striking feature of the table is the greater popularity of pattern recognition methods, as opposed to more mechanical statistical methods.

Table 11. Use of techniques by market

Currencies	DEM	JPY	GBP	Distribution
Chart Patterns	204	176	202	42.30%
Trendlines	196	218	233	47.00%
Stochastics	172	173	194	39.20%
Fibonacci Numbers	141	147	143	31.30%
Moving Averages	125	124	136	28.00%
Channels	61	100	45	15.00%
Relative Strength	45	54	52	11.00%
Candlestick Charts	46	50	34	9.40%
Elliott Wave	57	55	36	10.80%
Others	29	20	25	5.40%
Congestion Breakout	13	8	12	2.40%
Total weeks x analysts	459	459	458	

Stock Indices	DJIA	NIKKEI	FTSE	Distribution
Chart Patterns	220	206	149	46.10%
Trendlines	174	164	147	38.90%
Stochastics	183	188	113	38.80%
Fibonacci Numbers	110	146	89	27.60%
Moving Averages	140	141	78	28.80%
Channels	90	55	84	18.30%
Relative Strength	48	47	46	11.30%
Candlestick Charts	47	54	49	12.00%
Elliott Wave	33	28	23	6.70%
Others	38	40	22	8.00%
Congestion Breakout	23	22	20	5.20%
Total weeks x analysts	454	453	341	

Notes: Table shows number of times weeks x analysts in which each technique was mentioned. Techniques mentioned in fewer than 1% of months were Gann 0.8% , Cycles 0.3% , Seasonal 0.3% , Market Profile 0.5 and Volume/ Open Interest 0% . The only "Other" technique explicitly listed by respondents but not included in the questionnaire is the moving average convergence-divergence MACD indicator.

The most cited techniques are "Chart Patterns" and "Trendlines", with related methods such as "Channels" also prominent. "Fibonacci Numbers" were cited more frequently than "Moving Averages", and the exotic "Candlestick Charts" and the "Elliott Wave" appeared in about 10% of all the surveys. The most popular type of

statistical indicator is “Stochastics”, followed by the more intensively researched “Moving Averages” and “Relative Strength”. The use of these indicators differs little between currency forecasting and equity index forecasting. The only noticeable difference is the greater use of Candlestick Charts and Congestion Breakout in the equity markets. An immediate conclusion from the Table is that the emphasis in the academic literature on easily replicable statistical methods like moving averages is misplaced.

The use of indicators does differ across forecasters. Tables 12A and 12B show for currencies and stock indices respectively the methods of each analyst. The first column shows the number of techniques usually mentioned by the analyst. Most combine 2-3 methods together. Only forecaster I (a relatively accurate and profitable forecaster as noted above) switches between one method and another, and usually only mentions one technique. Again, this underlines the lack of realism in academic studies that focus on trades based only on one indicator.

The remaining columns of Tables 12A and 12B show the mix of techniques favoured by each analyst. Many forecasters appear to have a preferred mode of working, which they apply to both currency and equity forecasting. Forecaster J, for example, cites Chart Patterns, Trendlines and Moving Averages in nearly all weeks for nearly all targets. Forecasters B and F often use Stochastics. Forecaster A often mentions moving averages. Forecaster T relies on Elliott Wave analysis allied to Fibonacci Numbers. Rather curiously, forecaster I uses the Elliott Wave analysis not for stock market forecasting (where it was developed) but for currency forecasting, where there is less rationale for its use. Other forecasters like L and P are more eclectic, and shift emphasis from one indicator to another over time. The overall impression from the survey responses is that technical analysts are quite heterogeneous. They use different techniques. Some stick with their preferred methods regardless of market conditions. Others change the emphasis that they place on different indicators as the environment changes.

Table 12A. Use of Techniques by Forecaster: Currencies

Forecaster	Modal	% of weeks mentioned:										
	No. of techniques	Chart Patterns	Trendlines	Stochastics	Fibonacci Numbers	Moving Averages	Channels	Relative Strength	Candle- Sticks	Elliott Wave	Congestion Breakout	Others
A	3	28%	53%	0%	78%	89%	6%	0%	0%	0%	0%	20%
B	2	36%	54%	77%	5%	46%	1%	15%	1%	0%	0%	0%
F	3	16%	26%	96%	33%	0%	24%	1%	54%	0%	0%	1%
G	3	57%	66%	60%	48%	0%	37%	4%	0%	0%	0%	0%
H	2	4%	23%	59%	12%	72%	10%	1%	0%	0%	2%	1%
I	1	33%	2%	21%	2%	2%	2%	1%	2%	54%	2%	11%
J	3	96%	99%	0%	1%	97%	0%	4%	0%	0%	0%	0%
K	3	70%	49%	0%	46%	0%	33%	66%	11%	1%	7%	0%
L	3	31%	38%	60%	26%	8%	26%	38%	25%	0%	11%	27%
M	2	12%	27%	55%	57%	0%	58%	0%	7%	0%	5%	2%
N	2	79%	83%	0%	0%	0%	15%	0%	0%	0%	0%	3%
P	2	16%	67%	48%	57%	2%	0%	4%	17%	16%	5%	2%
S	3	97%	53%	54%	18%	47%	3%	6%	0%	3%	0%	0%
T	2	19%	28%	14%	72%	0%	0%	0%	0%	85%	1%	0%
All	2	42%	47%	39%	31%	28%	15%	11%	9%	11%	2%	5%

Table 12B. Use of Techniques by Forecaster: Stock Indices

Forecaster	Modal	% of weeks mentioned:										
	No. of techniques	Chart Patterns	Trendlines	Stochastics	Fibonacci Numbers	Moving Averages	Channels	Relative Strength	Candle-Sticks	Elliott Wave	Congestion Breakout	Others
A	3	70%	3%	4%	53%	92%	0%	0%	0%	0%	0%	42%
B	2	43%	49%	71%	2%	49%	0%	10%	0%	0%	0%	2%
F	3	14%	16%	95%	34%	2%	33%	3%	56%	0%	1%	5%
G	3	52%	67%	43%	62%	7%	43%	1%	1%	0%	4%	0%
H	2	5%	16%	65%	10%	70%	11%	4%	0%	0%	8%	0%
I	1	44%	5%	35%	3%	2%	2%	4%	5%	3%	4%	18%
J	3	95%	98%	1%	0%	95%	0%	1%	0%	0%	1%	0%
K	3	67%	32%	0%	41%	0%	50%	66%	17%	1%	8%	0%
L	3	36%	37%	64%	15%	4%	19%	35%	17%	0%	24%	35%
M	2	18%	17%	48%	58%	0%	50%	2%	7%	0%	10%	2%
N	2	59%	87%	0%	0%	2%	14%	0%	0%	0%	0%	0%
P	3	19%	38%	43%	48%	1%	14%	6%	49%	29%	4%	1%
S	3	100%	54%	49%	3%	83%	0%	3%	0%	2%	2%	0%
T	2	21%	21%	14%	78%	2%	6%	0%	0%	90%	0%	0%
All	2	46%	39%	39%	28%	29%	18%	11%	12%	7%	5%	8%

In principle, some techniques such as moving averages and trendlines are appropriate in trending markets, while others such as chart patterns and Stochastics are useful in situations where the market is consolidating or reversing direction. It is therefore interesting to ask whether analysts switch between techniques in a systematic way depending on market tone. Table 13 shows some weak support for this idea. We have measured the correlation between the mention of a trending or consolidating market tone, with mention of selected techniques. Use of moving averages and trendlines are indeed less at times when the market is perceived as consolidating. Chart patterns however are also used more intensively at such times, suggesting that all of the analysts tools (except possibly stochastics) are directed at trending markets.

Table 13. Correlation of technique and market tone

Techniques	Trending		Consolidation	
	Currencies	Equities	Currencies	Equities
Moving Averages	0.35	0.36	0.17	0.17
Stochastics	0.10	-0.16	-0.01	0.39
Trendlines	0.47	0.43	0.06	0.16
Chart Patterns	0.27	0.50	0.28	0.22

Notes: Table shows correlation between mention of trending or consolidating market tone, and mention of selected chart techniques.

Do some techniques have more predictive power than others? To test this we have pooled data from all weeks and forecasters, and regressed the (absolute) forecast error, and the profits from each forecast, on dummy variables for each major technique. The dummies take the value of 1 if the technique is mentioned by the forecaster that week, and 0 otherwise. The regressions have the form:

$$Y_{it} = a_i + b_i \sum_{j=1}^{11} D_{ijt} + u_{it} \quad (4)$$

where Y_{it} is the error in the forecast made by forecaster i in week t (or the profit made between week t and $t+1$ by trading on the forecast), and D_{ijt} takes the value 1 if forecaster i mentions technique j in the forecast made in week t .

Table 14A and 14B show the regression results for currencies and stock indices respectively. The overall fit of the regressions are very poor, suggesting very little of the variation in week-to-week and forecaster-to-forecaster performance can be attributed to differences in forecast method. There are a few statistically significant results – for example, use of Channels and Fibonacci ratios was associated with larger forecast error in predicting the Deutschmark-Dollar exchange rate – but these are no more than would be expected by chance in a data set with no underlying structure. Perhaps unsurprisingly given (a) the survival of our forecasters in a competitive marketplace, and (b) our earlier finding of few significant differences in forecast accuracy ranks, there are no obviously dominated or dominating forecast methods.

Table 14A. Regressions of forecast error and trading profits on forecast technique: currencies

Dependent variable - absolute % error						
Variable	DEM		JPY		GBP	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
C	0.0174	7.23	2.9209	8.51	0.0167	9.04
Moving Averages	-0.0001	-0.02	-0.3353	-1.26	-0.0003	-0.20
Stochastics	0.0009	0.51	-0.2528	-1.03	0.0001	0.06
Relative Strength	-0.0016	-0.60	-0.0413	-0.13	-0.0004	-0.19
Trendlines	-0.0005	-0.28	0.0897	0.39	-0.0004	-0.32
Channels	0.0066	2.62	0.0696	0.25	0.0041	1.92
Chart Patterns	0.0019	1.09	-0.4788	-1.99	-0.0005	-0.35
Elliott Wave	0.0008	0.26	-0.1775	-0.44	-0.0016	-0.66
Gann	0.0063	1.02	-2.3084	-1.43	0.0185	1.39
Fibonacci	0.0041	2.28	-0.2752	-1.18	0.0007	0.50
Candlesticks	-0.0007	-0.28	-0.5683	-1.65	-0.0022	-0.96
Congestion breakout	-0.0074	-1.72	-1.0382	-1.50	0.0039	1.14
R-squared	0.0387		0.0249		0.0205	

Dependent variable - % profits						
Variable	DEM		JPY		GBP	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
C	0.0056	1.79	1.0910	2.03	0.0028	1.09
Moving Averages	-0.0016	-0.62	0.2255	0.54	-0.0006	-0.27
Stochastics	0.0002	0.06	-0.1235	-0.32	0.0005	0.28
Relative Strength	0.0009	0.25	-0.2460	-0.49	0.0008	0.32
Trendlines	-0.0015	-0.68	-0.1574	-0.43	0.0009	0.53
Channels	-0.0030	-0.92	0.1683	0.38	0.0014	0.48
Chart Patterns	0.0001	0.04	0.1228	0.33	-0.0007	-0.38
Elliott Wave	-0.0013	-0.34	0.2307	0.37	-0.0024	-0.71
Gann	-0.0074	-0.92	1.0866	0.43	0.0377	2.06
Fibonacci	-0.0052	-2.19	0.1267	0.35	0.0017	0.87
Candlesticks	-0.0013	-0.40	0.3858	0.71	0.0010	0.30
Congestion breakout	0.0113	2.02	0.6460	0.60	-0.0036	-0.75
R-squared	0.0250		0.0048		0.0165	

Table 14B. Regressions of forecast error and trading profits on forecast technique: stock indices

Dependent variable - absolute % error						
Variable	DJIA		NIKKEI		FTSE	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
C	209.30	7.09	637.49	1.41	155.22	7.89
Moving Averages	-2.57	-0.11	-18.28	-1.16	-26.17	-1.40
Stochastics	-6.10	-0.28	-56.56	-0.91	-11.23	-0.69
Relative Strength	21.04	0.66	-40.96	-0.66	-14.61	-0.71
Trendlines	-18.76	-0.90	-64.88	0.91	36.19	2.34
Channels	5.61	0.21	34.88	0.99	-0.69	-0.04
Chart Patterns	13.45	0.62	-66.62	-0.41	-2.96	-0.19
Elliott Wave	-66.95	-1.63	-160.33	0.01	-7.53	-0.25
Gann	253.17	2.49	-411.00	0.25	-57.88	-1.11
Fibonacci	5.00	0.21	-55.16	0.76	9.38	0.55
Candlesticks	29.74	0.88	-41.21	0.00	48.63	2.37
Congestion breakout	-42.40	-0.95	303.34	-1.20	2.98	0.10
R-squared	0.03		0.05		0.05	

Dependent variable - % profits						
Variable	DJIA		NIKKEI		FTSE	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
C	42.92	1.25	111.78	1.41	-20.05	-0.87
Moving Averages	-9.54	-0.34	-74.37	-1.16	34.01	1.56
Stochastics	19.87	0.79	-56.63	-0.91	4.90	0.26
Relative Strength	21.00	0.56	-59.01	-0.66	35.46	1.47
Trendlines	-14.27	-0.59	54.50	0.91	-15.10	-0.83
Channels	-20.67	-0.67	83.77	0.99	-12.14	-0.61
Chart Patterns	15.65	0.62	-24.75	-0.41	-0.74	-0.04
Elliott Wave	15.05	0.31	1.76	0.01	31.86	0.90
Gann	5.93	0.05	145.43	0.25	247.80	4.05
Fibonacci	-0.77	-0.03	48.03	0.76	0.47	0.02
Candlesticks	4.03	0.10	-0.19	0.00	-20.13	-0.84
Congestion breakout	-3.42	-0.07	-153.16	-1.20	30.55	0.90
R-squared	0.01		0.02		0.07	

4.6 CONCLUSIONS

This study has explored the performance of a group of technical analysts, and the methods they used in preparing their forecasts. Although the track records of economic forecasters have been examined in the past, this study is to our knowledge the first to focus on individual technical analysts.

Our findings can be summarised into four propositions.

First, the analysts are very heterogeneous. The wide variety of technical tools discussed in specialist textbooks are all used to some degree by members of our panel. Some have preferred sets of tools. Others switch from one to another. All use more than one forecasting technique to support their trading recommendations.

Second, while these variations in technique translate into variations in point forecasts and profits across analysts, these differences in accuracy do not appear statistically significant. In similar vein we have found that there is no obvious correlation between the use of particular technical analysis techniques, and forecasts accuracy and profitability. So no one forecaster appears significantly better than the others, and no one technique appears significantly better either.

Third, differences across forecasters in accuracy as measured by mean absolute and mean squared errors are largely uncorrelated with differences in profits from trading on their forecasts. There is however, a weak positive relationship between directional accuracy and trading profits.

Finally, in this panel there is clear evidence of expertise in currency forecasting. Following most analysts would have resulted in excess risk-adjusted profits, relative to buy-and-hold FX benchmarks. Moreover, the Consensus forecast – the unweighted average of individual forecasts – gives reliable profits in excess of most individual forecasters. The same cannot be said of stock index forecasts, and most individuals

struggle to match buy-and-hold index benchmarks. Whether this reflects a difficulty in predicting stock indices relative to currencies, or whether this reflects the relative expertise of our panel, is harder to determine. Probably both factors apply.

CHAPTER 5

DAILY FOREIGN EXCHANGE COMMENTARIES VERSUS TRADING RULES

5.1 INTRODUCTION

The aim of this chapter is to compare the real-time trading recommendations of technical analysts in foreign exchange markets with the performance of mechanical trading rules based on the indicators that the analysts claim to use. We use a unique database which matches individual analysts' statements (made at the beginning of each day) about targets, trading ranges and support and resistance levels, with their recommended trading positions as exchange rates evolve through the day. We find that for the four rates studied – the US Dollar prices of the Euro, Sterling, Yen and Swiss Franc, in the years 2000-1 - the profits from the recommended trading positions are positive and significantly higher than the profits from following mechanical rules based on the same indicators. Many academic studies attempt to evaluate technical analysis by assuming that traders follow simple trading rules. Our findings suggest that this biases these studies against finding value in technical analysis.

Technical analysis is overwhelmingly the most popular method used by traders to interpret and forecast very short term movements in exchange rates. This has been well documented in the surveys of forecasting practices of foreign exchange dealers in London by Taylor and Allen (1992), in other European markets by Menkhoff (1997) and Oberlechner (2000), and in Hong Kong by Liu and Mole (1998). However, academics have struggled to establish whether technical analysis “works”. There is little direct evidence on the performance in practice of technical trading systems. Most studies instead generate a set of technical indicators *ex post* from a time series of exchange rate or stock price data, and examine the profitability of replicable trading rules – for example, filter rules, breaks from trading ranges, or moving average crossovers - based on these indicators. Sometimes this generates excess profits and

sometimes it does not, with success and failure depending more on the data set used than on the trading rule adopted.

The paper starts in Section 2 with a survey and critique of this growing body of research on technical analysis. Section 3 introduces our data, which consist of start-of-day commentary on the market for all trading days in the years 2000-1, a set of trades recommended during the day following the commentary, and a matched hourly time series of spot exchange rates. Section 4 examines the accuracy and calibration of the forecasts given in the morning commentaries. Section 5 compares the profitability of the recommended trades with the profitability of purely mechanical trades. These include simple filter rules, positional trades based on published exchange rate targets, and range-breaking rules based on support and resistance levels identified in the morning commentary.

We find that the analysts' forecasts have little directional accuracy, but that their recommended trades are highly profitable. We also find that the recommended trades are significantly more profitable (and less costly and less risky) than the trades based on mechanical rules. The distribution of time spent in the market for recommended trades also has a very different pattern from the most distributions produced by the mechanical rules. In Section 6 we show that the recommended trades may be related to the support and resistance levels, but the relationship is not simple. Moreover, once in place the trades are subject to very tight stop-loss limits and profit-taking limits, and as a result are skewed towards shorter holding periods than most mechanical rules. In practice the link between the technical indicators, recommended trades, and trading profits is more complicated than academics have assumed, and indeed more complicated than the rules enunciated in standard textbooks on technical trading.

5.2 STUDIES OF TECHNICAL ANALYSIS

Technical analysis is not a homogeneous body of knowledge. There are many different technical indicators, some easy to model, others highly subjective, and analysts do not necessarily use these singly, or in a consistent way. The range of technical indicators is well illustrated in standard industry texts such as Pring (1998), Murphy (2000) and Edwards and Magee (2001), and Neely (1997) provides a digestible summary with illustrations from the foreign exchange market.

Technical indicators fall into three classes. First, there are purely statistical indicators, which can be easily formalised and evaluated. Commercial technical trading software has for many years offered traders the facility to “backtest” statistical trading rules. An example is the moving average rule, which recommends that the trader is long or short depending on whether the current price, or a short term moving average of recent prices, is above or below a moving average of prices. The studies of stock market prices and indices by Brock et. al. (1992), and Gunasekarage and Power (2001) report positive excess profits from simple moving average rules, and from neural network systems with moving average inputs (Gencay, 1996, 1998). However, parallel studies by Lee and Mathur (1994), Lee, Gleason and Mathur (2001), Lee, Pau and Liu (2001) and Olson (2003) have failed to find consistently profitable moving average trading rules in currency markets.

The second class of trading rules rely on pattern recognition by analysts. This can mean either the identification of a turning point formation, or of a trading channel.

Academic studies of turning point formations are relatively recent, with researchers using computationally intensive data mining methods to identify and test the impact of turning point patterns. Examples are the studies of the “head-and-shoulders” pattern in stock prices and exchange rates by Osler (1998) and Chang and Osler (1999), and the variety of reversal patterns identified in stock prices by Lo, Mayansky and Wang (2000). These studies show that specific patterns identified by chart analysts do occur around turning points more often than would be expected by chance. But whether this can be translated into profits is debateable. Lo et. al. (2000)

note that the distribution of stock price movements changes after classic reversal patterns, but do not explore whether this could be used to trigger profitable trades. Chang and Osler (1999) find little evidence that the Head and Shoulders pattern can be exploited to make profits in major currency markets.

Academic studies of trading ranges have a longer pedigree. One series of studies – originally designed as tests of weak form market efficiency – has investigated the profitability of “filter rules”, whereby traders buy (sell) if the price rises (falls) by more than some critical amount. Typically, the upper filter is set at some fixed percentage (say 1%) above the low price encountered during the current position, and the lower filter is set the same percentage below the latest high price. Stock market examples include Alexander (1961), Fama and Blume (1966) and Sweeney (1988). Sweeney (1986) and Levich and Thomas (1993) test the profitability of filters of various sizes using time series of daily closing price data on major currencies in the years 1975-80 and 1976-1990 respectively. They find that small filters (0.5%-1%) yield profits significantly greater than would be expected by chance.

A filter rule is a poor caricature of how trading ranges are defined in practice. Murphy (2000) and Edwards and Magee (2001) do suggest that support and resistance might simply be at the most recent low and high price, which would resemble a simple filter rule. But more often support and resistance are determined by drawing trendlines through a longer series of past lows and highs. Support and resistance can also occur at psychologically important levels. For example, DeGrauwe and Decupere (1992) document the clustering of exchange rates close to “round numbers” – exchange rates ending in a 0 or a 5 – and similar barriers can be found in stock markets (Donaldson and Kim, 1993). It is now recognised that this reflects more than just market psychology. LeBaron (1996) and Szakrany and Mathur (1997) note that exchange rates often cluster around the limits of bands defended by central banks, as they intervene to offset the effects of market buying or selling. And Osler (2001) shows that retail stop loss and limit orders, and options exercise prices, often occur at round

numbers, so it is rational for traders to expect price movements to be arrested at such levels.

A few studies have examined technical indicators generated by technical analysts and made available to the market. Goodhart and Curcio (1992) and Curcio et. al (1997) gathered information from the Reuters information service on published support and resistance levels in the currency markets, and examined the profitability of plausible trading rules based on these indicators. The idea that traders act on published support and resistance levels like these is strongly supported by Osler (2000, 2001). She demonstrates that there is a high degree of agreement among analysts from different institutions about where support and resistance will be found, and that transactions-level data on exchange rates tend to cluster around published support and resistance levels.

As with the filter rules, trading results from these studies have been inconsistent. Using hourly data on four major currencies against the dollar from a 3-month period in 1989, Curcio and Goodhart (1992) found that a strategy of going long when the spot rate broke above the resistance level, and short when it fell through a support level, was profitable. However, when the experiment was repeated by Curcio et. al. (1997) on a 5-month stretch of data from 1994, the rules were not found to be profitable. The authors attribute the difference in performance to the different character of the markets in 1989 (trending) and 1994 (mostly non-trending).

These studies assume that traders use support and resistance levels only to identify breakouts from a trading range. This is not the only use of trendlines highlighted in technical analysis texts, however. If a price has fallen towards a support level but failed to break it, for example, this may be interpreted as a buy signal, on the expectation that the price will reverse towards the middle of the trading range. Similarly, failure to break resistance may be a reversal sell signal. Moreover, if a break through resistance or support does occur, a common dictum is that "the old

resistance becomes the new support”, so that positions might be reversed rather than closed if the price falls back into the initial trading range.

Beyond these standard statistical and pattern recognition methods lies a hinterland of diverse and exotic charting techniques. Many technical analysts use alternative visualisations of data such as candlestick charts and point-and-figure diagrams. Many believe in recurrent but ill-defined “waves” in prices, exemplified by the popular book on Elliott waves by Prechter and Frost (2000). Many analysts use “magic numbers” such as Fibonacci ratios to fix price targets, and incidentally support and resistance levels. Some analysts believe in the influence of astrological forces. Either because of the difficulty of formalising these rules, or possibly fearing for their reputations as scientists, academics have not studied these more esoteric technical trading tools.

All technical analysts stress the importance of “confirmation” – that trades should be triggered only if several indicators point in the same direction. Moving average rules are known to be ineffective in non-trending markets, and so would not be used in practice without first testing – formally or informally - for the presence of a trend. Given the results of the above studies on filter rules and channel trading, it would also seem necessary to use these in conjunction with some other indicator measuring the strength of the market trend. Neely et. al. (1997) and Allen and Karjalainen (1998) examine nonlinear methods of combining statistical indicators, but with results as ambivalent as those from the simpler rules. However, Neely and Weller (1999) do show that the profits from applying technical rules singly are consistently lower than the profits from a more complex genetic programming rule that effectively combines several indicators. Many analysts also use information on economic fundamentals together with the price data that underpin technical analysis. The significance of environmental information is emphasised in the recent surveys of traders’ practices and beliefs in the London and New York markets by Cheung et. al. (2000) and Cheung and Chinn (2000). While the academic studies surveyed in Karpoff (1987) have certainly examined the information content of trading volume, there has been

little work on the effects of combining technical data with broader economic fundamentals.

A criticism of all these exercises is that neither the trading rules, nor - in most cases - the indicators, were ever actually used by traders. The indicators studied are often statistical artefacts. They are almost always used singly, and no allowance is made for combining indicators. The prices used are often end-of-day quotes rather than real time transactions prices. Positions are opened and closed automatically, with no role for the exercise of judgment. A generous interpretation of the findings of the academic studies on technical analysis is that it would have been possible for traders to act on these rules, so any profits found might be regarded as the minimum attainable by more sophisticated technical traders. Even this has been challenged by some critics, on the grounds that the most carefully executed ex post studies such as that by Brock et. al (1992) are subject to data snooping biases (Sullivan, Timmerman and White, 1999).

5.3 DATA

Our data source is Standard and Poor's MMS (now MMS International). This service provides continuous commentary and analysis on financial markets, and the analysis is distributed worldwide to money market and foreign exchange dealing rooms via the screens of all major quote vendors, including Bloomberg, Reuters and Telerate, and through the company's website www.globalmarkets.com. The service has many thousands of subscribers, including all leading international banks, and for a number of years has been voted "best screen-based service" by FX Week magazine.

The organisation employs technical analysts in London, New York, Toronto, San Francisco and Singapore. The analysts provide two pieces of information used in our study - a beginning-of-day commentary on forex markets, and through-the-day trading recommendations based on their commentary. We have matched their

commentary and trading recommendations with high frequency time series on spot exchange rates, also from MMS International.

5.3.1 DAILY COMMENTARY

Shortly after the opening of each trading day in London a local analyst writes a comment on technical features of the market for each major currency. The idea is that this will inform trader behaviour as the day progresses. The commentary is updated through the morning in London, and if there has been a significant change in the market the support and resistance levels will be changed. Then as the world turns and other markets open, the commentary is taken up by the North American and Far Eastern analysts.

In an earlier paper (Batchelor and Kwan, 2001) and chapter 4 in this thesis, we report the results of a survey conducted among the MMS analysts, designed to establish the type of technical analysis tools these analysts use. A large number of techniques – statistical, pattern based and exotic - were reported, and all analysts used more than one style of indicator. The most popular were trendline (support and resistance) and reversal pattern methods. There was agreement across the analysts on the primacy of these methods, which were regarded as core technical tools that would be widely understood by clients.

In this study we look at four markets, for the US Dollar against the Euro (EUR), Sterling (GBP), the Swiss Franc (CHF) and the Yen (JPY), through the years 2000-2001. Here is an example of the daily commentary relating to the Cable (i.e. Dollar-Sterling) rate as it appeared on traders screens on the morning of 23 May 2001. At the time the comment was published the spot rate was 1.4218.

10:35 GMT - MMS LDN - [Cable] continued to
weaken in Far Eastern markets after
falling sharply below its recent trading

range of 1.4300-1.4500. The negative market tone calls for a target of 1.4160 but with some support down to 1.4180. Provided there is no break below this level, potential further out is for a recovery above 1.4240. [PA]
 [Target 1.4160 Res: 1.4237 1.4292 Sup:1.4186
 1.4154 Range 1.4300-1.4175]

This comment is typical of all the data used in our study. Following the verbal comments, the analyst [PA] gives four pieces of numeric information summarising his view of the market:

- a **target** for the exchange rate of 1.4160
- a **range** for the day of 1.4175-1.4300
- two **resistance levels** which we label $R1 = 1.4237$ and $R2 = 1.4292$
- two **support levels** labelled $S1 = 1.4186$, $S2 = 1.4154$

The target gives information on the expected direction of change of the rate, and the size of the expected change. The trading range gives an estimate of the expected high and low for the day. Consistent with the bearish tone of the forecast, it is centred below the current spot rate. The target in this case is outside the forecast daily trading range, with the implication that the analyst does not expect the target to be achieved within the day.

Resistance and support levels bracket the current spot rate, and as discussed above indicate points at which movements up or down in the rate are likely to be arrested by selling or buying pressure respectively. If we imagine that these levels are derived from local trendlines drawn through recent highs and lows for the price series, then $R1$ and $S1$ would come from a short window of recent data, while the broader levels $R2$ and $S2$ would reflect the trading range in a longer window of data. Note that the resistance and support levels need not be within the expected daily trading range.

On Table 1 we have set out some stylised facts about the support and resistance levels, and the trading recommendation levels, for Sterling. The pattern for other currencies is similar. The upper panel of the Table shows that in our data set the support and resistance levels tend to be symmetric around the spot rate at the time the commentary is written, with the R2 / S2 band about twice as wide as the R1/ S1 band. The lower panel of the Table shows that “rounding” is a prominent feature of our data. About 55-60% of the support and resistance levels have a fourth decimal digit rounded to either 0 or 5, and 20%-30% have the a third digit rounded to 0 or 5, as against an expected frequency of 10% if the digits were random. The proportions are even higher when we look at levels where the analysts set targets and stop loss limits (90-95% end in 0), and equally high when we look at levels where their recommended trades are opened or closed.

Table 1. Stylised facts on support, resistance, and trading levels, GBP 2000-1.


The upper panel of table of this Table shows difference in ticks between spot rate and support and resistance levels in daily commentary, and the incidence of the digits 5 and 0 as the trailing digits in the analysts' support and resistance levels. The lower panel shows the difference between the entry price of each recommended trade and the target, stop and close prices of the trade, for both short and long positions together. The table also shows the incidence of the digits 5 and 0 as the trailing digits in the analysts' claimed entry and close prices, and their target and stop-loss levels.

		3rd digit		4th digit	
		5	0	5	0
Commentary					
	Difference from Spot (ticks)				
R2	68	13.6%	19.1%	21.5%	41.1%
R1	32	12.7%	17.6%	20.2%	34.7%
Spot	0	9.7%	8.4%	9.0%	9.7%
S1	-29	15.8%	20.2%	20.9%	35.2%
S2	-65	9.2%	18.9%	20.2%	40.0%
Recommended Trades					
	Absolute Difference from Entry (ticks)				
Entry	0	17%	18%	31.9%	65%
Target	94	18%	31%	16.9%	78%
Stop	40	16%	14%	27.7%	71%
Close	47	16%	17%	31.2%	66%

Figure 1 shows end-hour bid prices for the spot dollar-sterling exchange rate through the 23 May 2001. Overlaid on the exchange rate data are the target, support and resistance levels for the previous day, the current day, and the following day. The new levels for 23 May from the commentary broadcast at 10:35 are plotted starting at 11:00 GMT, with the idea that only around this time could readers of the commentary act on the analyst's forecasts.

Figure 1. Support, resistance and recommended trades, GBP May 23-24 2001.

This Figure shows hourly rates for the US Dollar – Sterling exchange rate for the 24 hours from 06:00 on 23 May 2001 to 12:00 on 24 May 2001. Overlaid are support (S1 and S2) and resistance (R1 and R2) levels cited in the commentaries of 22-24 May, and the times/ prices during 23 May when analysts recommended long and short positions in Sterling.



5.3.2 TRADING RECOMMENDATIONS

Following the commentary an analyst may at some time during the day issue a buy or sell recommendation for the currency. During 23 May 2001, two trades were recommended, by different analysts, first PA, who wrote the commentary, and later SS who reversed PA's position. The recommendations were:

12:45 OPEN LONG at 1.4260 TARGET=1.4440 STOP-LOSS=1.4225 [PA]

15:50 CLOSE LONG at 1.4225 [PA]

15:52 OPEN SHORT at 1.4225 TARGET=1.4160 STOP-LOSS=1.4300 [SS]

20:45 CLOSE SHORT at 1.4185 [SS]

The trades are marked on Figure 1. Each trade consists of opening a position, and setting a target and a stop-loss level. Subsequently, the position is closed. The idea of setting a target is to give some indication of where profits might be taken on a winning position, and the idea of the stop-loss level is to set a price at which a loss-making position should be closed. In the cases above, the first position makes a loss of 35 ticks (1.4260-1.4225) and is closed out exactly at the stop-loss level. The second position makes a profit of 70 ticks, and is closed out at 1.4185, well before the target of 1.4160 is reached. In Section 6 below we show that this pattern of trading, with losses tightly controlled and profits taken ahead of target, is characteristic of the whole data set.

Note that the analysts do not give reasons for their trades. As it happens, the trades on 23 May can be rationalised in terms of the published support and resistance levels. The first long position is triggered by the break of price through the R1 resistance level. A common belief among analysts is that the old resistance level becomes the new support, so the switch into a short position might be triggered by the fall in price back through the R1 level. The short position is then closed at the S1 support level. In Section 6 below we examine whether, looking across all trades in all currencies, there are systematic relationships like this between the resistance and support levels from the commentary, and the subsequent recommended trading positions.

5.3.3 SPOT EXCHANGE RATES

The “actual” spot rates shown on Figure 1 start as high frequency data on bid quotes, also from MMS International. We use these for two purposes. First, we need notional

transactions prices at which mechanical trading rules are executed. For this purpose, we aggregate the data into end-hour prices. We require that these prices are quoted by at least two banks, so as to reduce the possibility of recording rogue quotes. This is the time series shown for a single day in Figure 1. In Figure 2 we show end-hour rates over the whole sample period 09:00 GMT 3 January 2000 – 07:00 GMT 20 December 2001. With allowances for holidays and other market closures, this amounts to just under 11000 hours of data. Figure 2 is scaled so as to reflect the value of an initial investment of \$1 in each foreign currency, and hence shows the returns in dollars to simple buy-and-hold FX strategies. The Figure ignores the effect of interest rate differentials, which would lead to a lower outcome for the long Yen strategy, but a slightly higher return for Sterling. The order of magnitude of these interest differential effects is small, and as we shall see has no bearing on the evaluation of the trading rules.

Figure 2 shows that our currencies experienced a mixture of trending and non-trending patterns. The Euro and Sterling were downward trended through 2000, but not in 2001. Overall they lost about 10% of their dollar value across the two years. The Swiss Franc did not show a trend in either year, but was subject to quite violent swings. The Yen was not trended in 2000, but depreciated sharply against the dollar through 2001, ending 20% lower.

Figure 2. Dollar exchange rates, hourly, 2000-1

Figure shows the “buy-and-hold” value of \$1 invested in each foreign currency on 3 January 2000, ignoring interest rate differentials.



The other use of our spot rate series is to audit the claims of the MMS analysts. Their “trades” are only recommendations and not actual trades, so the analysts do not have money at risk, only their reputation. This has the advantage that the analysts can be objective about the market to a greater degree than bank-based analysts, who are liable to be biased by their own trading positions. But the fact that the analysts are not implementing their trades raises three important questions. First, the prices at which trades are opened and closed are simply reported by the analysts, based on screen quotes - were any actual trades made at these prices? The predominance of round numbers in the trading recommendations suggests very strongly that the claimed entry and exit prices are indicative rather than real. To check this, we have matched the analyst prices against the original tick data, and all do appear genuine.

Second, could a user of the MMS service actually have traded at these prices, having first observed the analyst recommendation? This is more problematical. The delay in receiving and reacting to the recommendation means that any analyst prices will be stale by the time any trade is executed. In a parallel study of bond futures market trades (Batchelor and Kwan, 2003), we found that some analysts recommended trades at favourable but transitory prices, so the delay in trading may lead to a systematic overstatement of potential profits.

To test for and remove biases due to slippage and analyst optimism, we recompute their profits by assuming that recommended positions are entered only at the very end of the hour in which the recommendations are made. This puts the analysts’ trades on the same basis as the benchmarking mechanical rules described in Section 5 below. Where an analysts position is opened and closed within the same hour, we therefore assume a zero profit from the trade. In the case of the two trades in our illustrative example from 23 May 2001, the first position would have been opened at 13:00 hours at a price of 1.4240, and closed at 16:00 coincidentally at the same price, for a zero profit. The second trade would have been opened at 16:00 at 1.4240 and closed at 21:00 at 1.4190 for a profit of 50 ticks. As it happens, in this case the delay would have increased rather than decreased reported profits.

Finally, what about transactions costs, and the possible adverse impact on price from the act of buying or selling? The MMS forecasts are used by professional traders and all the currencies studied are very liquid, so spreads on each transaction are typically small, around 5 ticks for the example here. A spread of 10 ticks would be unusually high. However, both analysts and mechanical rules generate many trades, and – unlike the interest differential – transactions costs do materially affect profitability and must be factored into our calculations. When we compare actual and synthetic trading rules we first report profits unadjusted for transactions costs, but then consider their sensitivity to the bid-ask spread.

5.4 FORECAST ACCURACY AND CALIBRATION

The purpose of technical analysis is to provide trader support, not to provide point forecasts of future exchange rates, and the morning commentary does not contain an explicit forecast of this kind. However, it is interesting for two reasons to ask whether the commentary displays any forecasting expertise. First, it offers an opportunity to test on a large data set whether there is any relationship in our data between forecast accuracy and the profitability of trades based on the forecasts. Previous studies by Boothe and Glassman (1987) and Leitch and Tanner (1991) have found low or weak correlations, from very small samples of exchange rate and interest rate forecasts. Second, the commentary does contain an explicit forecast of the expected trading range, and it is possible to test whether the analysts are well calibrated in the sense that their ranges widen ahead of increased market volatility.

We use two measures of the expected exchange rate, based on analyst data. One is the target cited in the daily commentary. This is not really a forecast for a point in time, but rather an expression of where the rate might go over the next day or two if the expected trend materialises. The other measure is the mid-point of the high-low range for the day. If prices evolved as a random walk with drift, this would be an estimate of

the end-of-day expected exchange rate. But since the analysts usually have a more complicated view of how rate might evolve we should again treat this “forecast” with some caution. The mid-range measure is available for more days than the target. Of the 455 trading days in our sample period, we have commentary on Sterling for 441 trading days, and for all of these a trading range is given. However, on 30 of these days analysts do not give a quantitative target for the price, but make only a neutral verbal comment, such as “range trade”, “await reaction” or “flat position”, and on 47 other days no target is given at all. So the number of days on which we have quantitative information on both a trading range and a price target for Sterling is $441 - 30 - 47 = 364$. We use these days in our tests of accuracy, so that results can be compared across the alternative forecasts.

The upper part of Table 2 shows the exchange rates at the beginning and end of our sample, and their highs and lows for the years 2000-1. To test forecast accuracy, we need a measure of the daily change in the exchange rate. We use the average of the end-hour prices 23, 24 and 25 hours after each daily comment is published, with the idea that this will reduce any distortion due to the occurrence of outlying prices at the end of each 24-hour period.

We assess forecast accuracy in a conventional way, by computing the bias, root mean square error, and directional accuracy of the forecasts, and these error metrics are benchmarked against naïve alternatives. The lower panel of Table 2 summarises results for bias (average error, actual minus forecast) and root mean square errors in the analysts’ daily forecasts. Note that the Euro and Sterling rates are quoted here as US Dollars per unit of foreign currency, whereas the Swiss Franc and Yen are quoted as units of the foreign currency per US Dollar. Negative signs on the bias figures for the Euro and sterling, or positive signs for the Swiss Franc and Yen, means that forecasts did not fully anticipate the weakening in currencies against the dollar over the period of the study. The units of the bias and root mean square error figures are “ticks”. Again following market convention, the tick size for the first three currencies

is a 1 point move in the fourth decimal place of the quoted rate, but for the Yen it is a 1 point move in the second decimal place.

Table 2. Point accuracy of forecasts from daily commentaries, 2000-1.

The upper panel of this Table reports summary statistics on spot rates in the years 2000-1. Lower panel reports error metrics for analysts daily spot rate forecasts based alternately on their target, and the mid-point of their daily high-low range. Units of the exchange rate, bias and root mean square errors are ticks, defined as 1/10000 of the quoted rates for EUR, GBP and CHF, and 1/100 of the quoted rate for JPY.

		EURO	GBP	CHF	YEN
		\$ per €	\$ per £	per \$	per \$
<i>Number of days</i>		455	455	455	455
Exchange rates:					
	Jan-00	1.0114	1.6171	1.5852	101.74
	Dec-01	0.9016	1.4591	1.6320	128.65
	Low	0.8232	1.3715	1.5440	101.44
	High	1.0400	1.6568	1.8284	128.71
Daily changes (%)					
	Mean	-0.03	-0.02	0.01	0.05
	SD	0.77	0.53	0.74	0.61
Daily changes (ticks)					
	Mean	-2	-4	1	6
	SD	70	78	125	70
<i>Number of forecasts:</i>		359	364	344	359
Bias (ticks)	<i>Target</i>	0	-21	13	-7
	<i>Mid-range</i>	-7	-7	8	6
	<i>No-change</i>	-6	-3	1	5
RMSE (ticks)	<i>Target</i>	103	111	171	99
	<i>Mid-range</i>	76	81	145	70
	<i>No-change</i>	68	79	138	69
RMSE%	<i>Mid-range</i>	0.84	0.55	0.86	0.61

The ranking of the root mean square errors shows that the mid-range forecast is a consistently better predictor of the 24-hour-ahead price than the “target” rate. In percentage terms the daily forecast error is lowest for Sterling, and highest for the Swiss Franc, and all are under 1% per day. Both target and mid-price analyst forecasts are less accurate than the no-change random walk forecasts, though the difference in RMSE accuracy of the random walk against the mid-price forecasts is very small for all currencies.

Table 3 reports tests for directional accuracy of the mid-range forecasts based on 2 x 2 contingency tables comparing actual and forecast movements UP and DOWN in the exchange rates. Overall directional accuracy as measured by the fraction of correctly signed forecasts is poor, with the Euro and the Yen around .5, Sterling slightly below 0.5, and only the Swiss Franc perceptibly higher, at 0.58. The $\chi^2(1)$ statistics in Table 3 test the observed frequencies in the contingency tables against the frequencies expected in a no-skill benchmark. Under the no-skill null, the number of UP and DOWN forecasts is the same as analyst forecasts, but their incidence is independent of the actual UP and DOWN movements. So the upper left cell for the no-skill benchmark GBP forecast would be $173 \times 206 / 364 = 97.91$, and the overall directional accuracy 0.5. The statistics confirm that the directional accuracy of the mid-range forecasts is very poor, and only in the case of the Swiss Franc can the no-skill hypothesis be rejected.

Table 3. Directional accuracy of mid-range forecasts from daily commentaries, 2000-1

This Table shows the numbers of correctly signed and incorrectly signed analyst forecasts of daily spot rate movements in the years 2000-1. Directional accuracy is the fraction of correctly signed forecasts. The $\chi^2(1)$ p-value should be less than 0.05 for the observed directional accuracy to be significantly different from the no-skill benchmark at the 5% significance level.

EUR				GBP			
	Actual			Actual			
	UP	DOWN	Total	UP	DOWN	Total	
Forecast							
UP	78	94	172	97	109	206	
DOWN	83	104	187	76	82	158	
Total	161	198	359	173	191	364	
Directional accuracy:			0.51				0.49
Chi-squared (1) statistic			0.03				0.04
Chi-squared (1) p-value			0.85				0.85
CHF				JPY			
	Actual			Actual			
	UP	DOWN	Total	UP	DOWN	Total	
Forecast							
UP	109	61	170	95	78	173	
DOWN	83	91	174	102	84	186	
Total	192	152	344	197	162	359	
Directional accuracy:			0.58				0.50
Chi-squared (1) statistic			9.40				0.00
Chi-squared (1) p-value			0.00				0.99

Table 4 reports some statistics relevant to the calibration of the analysts – that is, their ability to predict the range of exchange rate movements. This is of interest because analysts can only define a channel within which a price can be regarded as trending, if they can to some extent forecast volatility. The first set of statistics on the table compares the average intra-day volatility (defined as the standard deviation of 24 end-hour prices) and average daily range of each currency, with the average daily forecast range. On average, the order of magnitude of the forecasts is very close to the actual high-low ranges. However, the second set of statistics show that forecast ranges are only weakly positively correlated with the previous day's range, and even more weakly correlated with the following day's volatility and trading range. There is

therefore some information in the range estimates relevant to predicting volatility - if the forecast range is narrow, it is more likely than not that volatility will be low. But because the directional accuracy of the mid-range forecasts is poor, there are many violations of the expected bounds, and currencies move outside the analysts' forecast high-low ranges on 80-90% of days.

Table 4. Calibration of analyst forecasts, 2000-1

This Table reports summary statistics on daily volatility of spot exchange rates. Intraday SD is the standard deviation of 24 end-hour bid prices, and daily range is the difference between the daily maximum price (Max) and minimum price (Min). Forecast range is the difference between the forecast High and forecast Low rate forecast in the analysts' daily commentary.

Currency:	EUR	GBP	CHF	JPY
Number of days:	431	441	433	431
Means (ticks):				
Intraday SD	25	27	45	25
Actual Daily Range	87	97	154	88
Forecast Daily Range	89	96	130	93
Correlation of forecast range with:				
Previous day range	0.21	0.22	0.24	0.26
Current intraday SD	0.11	0.15	0.06	0.15
Current day range	0.15	0.16	0.07	0.16
Violations (no. of days)				
Max > forecast High	156	156	220	188
Min < forecast Low	195	195	181	155
Total	351	351	401	343

5.5 PROFITABILITY OF RECOMMENDED AND MECHANICAL TRADES

The profitability of trading rules based on directional forecasts is known to be only weakly correlated with directional accuracy, even when the rules are simple and mechanical (Acar, 1998). This happens because changing volatility makes the distribution of daily exchange rate changes fat-tailed, and puts a premium on the ability to make correct directional forecasts in high volatility periods. In the case of our analysts' trading recommendations, which are presumably based on a non-mechanical combination of technical indicators with other information, there is even less reason to expect that the poor directional forecasting performance documented above will necessarily translate into a poor profit performance.

Table 5 reports profits from investing a notional \$1 in long or short foreign currency positions according to the analysts recommendations. In all cases, the apparent compounded profit at the end of the 2-year period is substantial, around 60% for the Euro and Swiss Franc, and around 80% for Sterling and the Yen. More profit was generated in year 2000 than in 2001 for all currencies. So the incidence of profits does not seem to be related to whether the currency was trending or not.

Because these are only claimed profits, they need to be tested for robustness to slippage and trading costs. As described above we have allowed for slippage by assuming that the trades take place not at the claimed prices, but at the prices at the end of the hour in which the trading recommendations are made. This increases the value of some trades, but more often reduces profits. Slippage has a very small effect (about -2%) on the profits from trading Yen, a slightly larger (-5%) effect on Sterling profits, and an appreciable effect on the Euro and Swiss Franc (-12%).

The effects of transactions costs are much greater, not suprisingly in view of the 350-450 transactions recommended for each currency. In Table 5 we report the value of the trading position adjusted for conventional market spreads of 5 and 10 ticks and – because these tick sizes are not fixed percentages of the underlying price –

alternatively for a uniform cost of .05% per round trip. The .05% spread reduces profits on the Euro and Swiss Franc by about 15%, and on the more heavily traded Sterling and Yen by around 20%. Taking slippage and transactions costs together, net profits over the years 2000-1 remain positive, at around 20% for the Euro and Swiss France, and 40% for Sterling and the Yen. A doubling in transactions costs to 0.1% would, however, eliminate profits on both the Euro and Swiss Franc.

Table 5. Profitability of recommended trades, 2000-1

This Table shows profits from recommended trades, measured as the value at end-2001 of \$1 invested at the start of year 2000. Interest earned when out of the market is not included in the calculation. Slippage is added by assuming that trades occur at end-hour prices.

Currency:		EUR	GBP	CHF	JPY
Number of trades					
	Long FX	178	260	214	187
	Short FX	173	195	180	304
	Total	351	455	394	491
Apparent Profit (Value of \$1)					
	Whole period	1.5969	1.7927	1.6391	1.8074
	2000	1.3093	1.4716	1.4079	1.4476
	2001	1.2196	1.2182	1.1642	1.2485
Adjusted for slippage		1.3973	1.6865	1.4616	1.7782
Adjusted for transactions costs					
	5 ticks	1.3303	1.5542	1.4642	1.4710
	10 ticks	1.1081	1.3474	1.3079	1.1971
	0.05%	1.3425	1.4472	1.3683	1.4262
Slippage x transactions costs		1.1747	1.3615	1.2202	1.4032

We benchmark these trading profits against profits from 6 mechanical rules. The rules are:

1. R1/S1 : LONG if Spot > R1, SHORT if Spot < S1, CLOSE when New Comment
2. R2/S2 : LONG if Spot > R2, SHORT if Spot < S2, CLOSE when New Comment

3. Hi/Lo: LONG if $\text{Spot} > \text{HIGH}$, SHORT if $\text{Spot} < \text{LOW}$,
CLOSE when New Comment
4. R1/S1 : LONG if $\text{Spot} > \max(\text{R1}, \text{HIGH})$, SHORT if $\text{Spot} < \min(\text{S1}, \text{LOW})$,
CLOSE when New Comment
5. Filter: LONG if $\text{Spot} > \text{MIN} + \text{FILTER}\%$, SHORT if $\text{Spot} < \text{MAX} - \text{FILTER}\%$,
CLOSE when New Signal
6. Target: LONG if $\text{Spot} > \text{TARGET}$, SHORT if $\text{Spot} < \text{TARGET}$,
CLOSE when New Comment

The first three rules are similar to the range breakout rules investigated by Curcio and Goodhart (1992), and Curcio et. al (1997), and discussed above. The data they use are average nearby support and resistance levels (equivalent to our S1 and R1), and average expected HIGH and LOW prices quoted by a number of analysts. Our second rule has the same structure as the first, but with trading only on breakouts from the broader R2/S2 channel. Curcio and Goodhart (1992) also consider “soft” range breaking rules similar to those advocated by Sweeney (1986, 1988), in which the price has to break by more than some critical amount before a position is taken. This makes no appreciable difference to their results or ours, and we do not report their performance here.

The Filter rule requires the trader to take a long position when the spot rate rises more than some percentage FILTER% from a recent low (MIN), and a short position if the spot rate falls more than FILTER% from a recent high (MAX). This leaves to be determined the questions of how far back we look to determine the MIN and MAX points, and how to fix the size of the FILTER%. The highs and lows we recompute for each new position – that is, we would go long if the spot price rose more than FILTER% from its low in an already short position, and vice versa. The filter levels thus resemble support and resistance lines, and change as new highs and lows are achieved, and as new positions are established. Most studies investigate a number of filter sizes, but that is not our aim here. To set the toughest possible benchmark for the analyst recommendations, we choose the FILTER% *ex post* so as to maximise profits

over our data set. This sets an upper bound to the profits that could be achieved by a fixed filter rule in real time.

Finally, the Target rule simply imagines that a trader goes long if the TARGET in the analyst's morning commentary is above the current spot, and vice versa, closing the position only when a new target is published.

Both the filter rule and the target rule mean that the trader will always be in the FX market, either long or short, and is never out of the market. This is not necessarily the case with the support and resistance type rules. A position will be opened by the break of a critical level, but closed - and not reversed - if either the spot price falls back through that level, or if the next day's commentary puts the spot rate in the middle of a new trading channel. The upper panel of Table 6 shows for Sterling trades the proportion of time in hours spent in long and short positions, and the time spent out of the market, under the different mechanical rules. Tables for the other currencies tell a similar story.

For the Filter rule and the Target rule, the trader is 100% in the market. The R1/S1 rule has the trader in the market about 60% of the time. For the Hi/Lo and R1/S1+Hi/Lo rules the proportion is about 40%, and for the R2/S2 rule the trader it is only 30%. The recommended trades also put the trader into the market only about 30% of the time. So both the R1/S1, and to a much greater extent the target and filter rules, exaggerate time spent in open FX positions.

The lower panel of Table 6 highlights another striking difference between actual and synthetic trading rules. About one third of the analyst recommended trading positions are closed within 1 hour of opening, and well over half are close within 6 hours. Only 4% of positions are held for more than 24 hours. This pattern is also observed in the range-breaking rules. In contrast, the filter rule has only 14% of positions closed within the hour, and nearly 60% held open overnight (and indeed over many nights).

The target-based rule by definition holds positions open until the next commentary, in around 24 hours time.

Table 6. Time in market under alternative trading rules, GBP, 2000-1

This Table shows the percentage of time spent in and out of the FX market, and the time distribution of active long and short market positions.

Trading Rule:	R1/S1	R2/S2	Hi/Lo	R1/S1/Hi/Lo	Filter	Target	Actual
Positions taken							
(number)							
Long FX	463	257	342	337	8	76	260
Short FX	334	220	259	255	5	74	195
Total	797	477	601	592	13	150	455
Length positions held							
(hours, % distribution)							
Long FX	25%	14%	19%	18%	23%	63%	18%
Short FX	33%	17%	24%	23%	77%	37%	13%
Out	42%	69%	58%	59%	0%	0%	70%
< 1 hour	31%	35%	30%	31%	14%	0%	34%
1-6 hours	31%	30%	32%	32%	21%	0%	22%
6-12 hours	11%	9%	10%	11%	0%	0%	22%
12-24 hours	24%	24%	25%	24%	7%	28%	18%
>24 hours	3%	2%	2%	2%	57%	70%	4%

Apart from these differences in the time distribution of positions of mechanical and recommended trades, there are also significant differences in the positions themselves. Table 7 shows the probability that, given that the analysts have recommended a position (short or long), the mechanical rule yields the same position. The correlation is greatest in the case of the Target rule. Days when analysts recommend long (short) positions are often days when the morning commentary forecast a rise (fall) in the rate. There is 65% agreement in the case of Sterling, around 70% for the Swiss Franc and Yen, and 80% for the Euro. But the correlation is much weaker for the range-based rules. For example, in the case of the R1/S1 rule for Sterling, the rule agrees with the analyst recommendation about 34% of the time, and takes the opposite position 24% of the time. The correlation is lowest for the Filter rule. Indeed, in the

case of three of the four currencies it is more likely than not that the Filter rule will take a position opposite to the analyst. For Sterling, for example, the probability of the Filter rule producing the opposite position is 57%.

Table 7. Conditional probability that mechanical trading position = recommended trading position.

This Table shows the fraction of hours when the trading position of each mechanical rule is the same as that of the recommended trade, conditional on the recommended trade being in the market (either short or long).

	EUR		GBP		CHF		JPY	
	Same	Opposite	Same	Opposite	Same	Opposite	Same	Opposite
R1/S1	0.38	0.21	0.34	0.24	0.31	0.25	0.31	0.21
R2/S2	0.18	0.10	0.20	0.11	0.17	0.11	0.16	0.08
Hi/Lo	0.25	0.18	0.25	0.17	0.22	0.20	0.24	0.15
R1/S1/Hi/Lo	0.24	0.17	0.24	0.15	0.21	0.19	0.23	0.14
Filter	0.39	0.60	0.42	0.57	0.40	0.60	0.62	0.38
Target	0.80	0.20	0.65	0.34	0.72	0.28	0.70	0.30

Given these differences in the number, duration and direction of trading positions, we should not be surprised to find differences in trading performance between mechanical rules and recommended trades, and among the mechanical rules themselves. The upper part of Table 8 shows profits from each rule, assuming trading at end-hour prices (so adjusted for slippage), but unadjusted for transactions costs. We look first at the range-breaking rules, then at the Filter and Target rules.

Only in the case of the Euro does the R1/S1 breakout rule generate significant profits. At 23% over two years, they are still only half of what the analyst recommendations produce. For Sterling and the Swiss Franc, the R1/ S1 rule generates no profits. For the Yen, the R1/ S1 rule results in a 16% loss relative to buying and holding dollars. Moreover, the R1/ S1 rule involves very frequent trading. The lower panel of Table 8 shows that it involved taking more than twice the number of positions in the Euro

recommended by the analysts, 731 trades as against 351 recommended, and so would be subject to much larger transactions costs. The broader R2/ S2 channel rule is more economical, but performs only marginally better than the R1/S1 rule for Sterling, the Swiss Franc and the Yen, and worse for the Euro. The Hi/Lo and combined Hi/Lo and R1/S1 range breakout rules offer no significant improvement.

Table 8. Profitability of mechanical trades

This Table shows profits from trades based on mechanical trading rules, measured as the value at end-2001 of \$1 invested at the start of year 2000. Interest earned when out of the market is not included in the calculation. All trades are assumed to be executed at end-hour prices.

Currency:	EUR	GBP	CHF	JPY
Profits: Compound value of \$1				
Trading Rule:				
R1/S1	1.2329	1.0333	0.9979	0.8402
R2/S2	1.0295	1.0144	1.0401	0.9218
Hi/Lo	1.1055	1.0302	0.9926	0.8850
R1/S1 + Hi/Lo	1.0662	1.0150	0.8587	0.9122
Filter	1.1204	1.0076	1.1029	1.2975
Target	1.0026	1.0957	0.9887	1.1747
Buy and Hold FX	0.8896	0.8958	1.0288	0.7908
Analyst Recommended	1.3973	1.6865	1.4616	1.7782
Number of trades:				
Trading Rule:				
R1/S1	736	797	804	800
R2/S2	393	477	515	388
Hi/Lo	571	601	649	545
R1/S1 + Hi/Lo	561	592	650	529
Filter	15	13	51	86
Target	140	150	154	133
Buy and Hold FX	0	0	0	0
Analyst Recommended	351	455	394	491


Our conclusion is that, contrary to the findings of Curcio and Goodhardt (1992), mechanical rules based on breaks through published support and resistance levels cannot generate reliable profits. Contrary also to the conjecture in the follow-up study by Curcio et. al. (1997), the success and failure of these rules also seems unrelated to whether or not the spot rate is trending. Moreover, the profits generated from these rules give no indication of the relative or absolute levels of profit which traders can produce, based on the technical indicators. The R1/S1 rule profits are highest for the Euro and Lowest for the Yen, but the recommended trades produced most profit from the Yen, and least from the Euro.

The optimised Filter rule gives positive profits for all currencies, with most profit coming from the Yen. The optimum filter sizes are around 2% for the Swiss Franc, 3% for Sterling, and 4% for the Euro and the Yen, much larger than those found in the earlier studies of Sweeney (1986) and Levich and Thomas (1993). As is evident in Tables 6 and 8, our large filter sizes lead to a much smaller number of trades than recommended by the analysts – only 13 positions over two years in the case of Sterling. On the one hand, this does lower transactions costs, though not by enough to make the profits from the filter rule competitive with the recommended trades. On the other hand, it means that the filter rules are not in any way mirroring the trading behaviour of the analysts. Smaller filter sizes would lead to more frequent trading. However, in our high-frequency data set profitability is very sensitive to filter size, much more than with the daily closing price data used in earlier filter rule studies. Any decrease in filter size aimed at increasing trading frequency towards, say, 50 positions per year, quickly leads to losses in all currencies.

The Target-based trading rule does not produce significant profits for any currency except the Yen, and in the case of the Swiss Franc makes a small loss. Recall that Table 3 showed that directional accuracy was only 50% for the Yen, but significantly better than (58%) for the Swiss Franc. It appears that across our currencies there is no association between directional forecasting accuracy, and the profitability of taking daily positions based on the directional forecasts.

It is natural to ask whether the higher profits of the recommended trades are being bought at some increased risk. However, the sheer scale of the difference in profits makes it unlikely that a risk premium could be the explanations for the analysts' success. The time pattern of analyst profits also make it unlikely that risk is a relevant factor. Figure 3 shows how the position based on recommended Sterling trades evolved through the years 2000 and 2001, alongside the value of the R1/S1, Filter rule and Target rule positions. We also show the paths for Euro trades, where the mechanical rules make the best showing. In both currencies, the recommended trading position is both more profitable and less risky, with lower volatility and a smaller maximum drawdown than the other strategies. The same pattern is observed for the Swiss franc and Yen.

Figure 3. Value of positions based on recommended trades and selected mechanical rules, GBP and EUR, hourly, 2000-1.
(Value of \$1)



5.6 THE LOGIC OF RECOMMENDED TRADES


The range breaking rules like R1/S1 and R2/S2 do a reasonable job of mimicking the distribution of trades by position and duration, but fail to yield any consistent profits. In this final section of the paper we explore two possible reasons why these mechanical rules have failed. First, while breaks through support and resistance are important, entry to and exit from positions are in practice triggered by many patterns in prices around these levels. Second, once an analyst has recommended a position, the position is actively managed by imposing stop loss and price targets, which may have little to do with earlier support and resistance levels.

Figure 4 sets out the pattern of entry prices and exit prices for short and long recommended positions in Sterling, relative to the spot, support and resistance levels current at the time the trades were recommended to be executed. For each trade we have found the critical level closest to the entry price and the price at which the trade was closed. For trades opened or closed outside the extreme R2 and S2 bounds we distinguish those which occur close to these levels (labeled "Above R2" and "Below S2"), and those which occur well outside ("Higher" or "Lower") the R2/S2 trading range. "Close" to R2 and S2 means less than $(R2-R1)/2$ above R2, and less than $(S1-S2)/2$ below S2, making the ranges above and below R2 and S2 symmetrical. Similarly, trades within the (R2, S2) range are assigned to the nearest support or resistance level. So a position opened at a rate between the opening Spot and R1, say would be described as "Above Spot" if it were nearer to the spot rate, or "Below R1" if it were nearer to R1. On this basis, Figure 4 shows the percentages of long and short positions inside the trading range that were opened above and below the critical levels R2, R1, Spot, S1, and S2.

If the R1/S1 breakout rule was popular in practice, we should find a bunching of entries into long positions just above the R1 level, and a bunching of entries into short positions just below the S1 level. There is some evidence for this in the figure. It is more common to enter long positions just above R1 than just below, and more

common to enter short positions just below S2 than just above. If the R2/S2 breakout rule is used we should see a bunching of entries into long positions above the R2 level, and entries into short positions below S2. Again, there is some visual evidence for this kind of behaviour in the sense that there are more entries into long positions above R2 than in the symmetrical range above. Similarly, there are more entries into short positions below S2 than above. Taken together, selling on breaks through S1 or S2 account for about 34% of all short position taken within the (S2, R2) range, and buying on breaks through R1 or R2 account for about 24% of all long positions.

Figure 4. Entry points for Long and Short Positions, relative to Support and Resistance Levels: GBP



This does support the idea that trading recommendations are to some extent based on simple breakout rules. However, most trades in Sterling were not of this kind. And the evidence from other currencies is less compelling.

Looking again at Figure 4, it seems that many long positions are entered just above the S1 and S2 support levels. As noted in our review of trendline-based trading rules, this could be interpreted as reversal trade, based on a prediction that the exchange rate will return to a higher point within the R1/S1 or R2/S2 trading channel. Or it could reflect the "old support becoming the new resistance", with traders buying after the price has dipped below then broken back above the S1 and S2 levels.

Around 40% of trades in Sterling take place at prices well outside the R2/S2 range. We have already noted that prices often drift outside the High/Low band expected by the analysts in their morning commentary. The R1 and S1 levels are on average around 30 ticks away from the spot rate, and the R2 and S2 levels about 70 ticks away, so that the rate has to change by only about 0.5% in the day to break the R2/S2 range. The many trades that are recommended well outside the R2/S2 band suggest that very often analysts are using triggers other than their publicly announced support and resistance levels.

Table 9 shows the distribution of entry prices for all currencies. To test formally for the presence of breakout rules we have tested hypotheses of the form

$H_0: (\% \text{ above level} - \% \text{ below level}) = 0$

against

$H_1: (\% \text{ long above level} - \% \text{ long below level}) > 0$ if level = R1, R2

$H_1: (\% \text{ short above level} - \% \text{ short below level}) < 0$ if level = S1, S2

We also test for the presence of reversal trading (buying if a support level is not broken, or selling if a resistance is not broken) by testing:

$H_0: (\% \text{ above level} - \% \text{ below level}) = 0$

against

$H_1: (\% \text{ long above level} - \% \text{ long below level}) > 0$ if level = S1, S2

$H_1: (\% \text{ short above level} - \% \text{ short below level}) < 0$ if level = R1, R2

The test statistics shown in Table 9 follow a standard normal distribution under the null, and we have marked in bold cases where the null is rejected at the 95% significance level. In the case of Sterling, there is significant evidence of breakout trading when the rate falls below S1 and below S2. However, the differences between the proportion of long positions in Sterling entered above and below R1 and R2 are not statistically significant. In other currencies, evidence of breakout trading is even weaker. For the Euro, the only significant feature is that more long positions are taken below spot than above the spot rate. The opposite is true of the Swiss franc, though given that it is quoted per dollar, this would mean both Euro and Swiss Franc were often bought on dollar weakness. For the Yen, the only significant result is that more long positions are taken below S1 than above, a pattern that cannot be rationalised in terms of breakout or reversal trading.

Table 9. Distribution of entry prices relative to support and resistance levels

The upper panel of this Table shows where recommended long and short trading position were opened relative to the analysts stated support and resistance levels. The middle panel reports test statistics for differences in the proportions of positions opened above and below each support and resistance level. These are normally distributed under the null hypothesis that the proportions are equal, and significant differences at the 95% level are shown in bold. The lower panel of the Table shows the percentages of trades that were opened outside the trading range expected at the beginning of the day.

	Open Short					Open Long				
	EUR	GBP	CHF	JPY	Total	EUR	GBP	CHF	JPY	Total
% distribution in range:										
Above R2	3%	5%	6%	4%	5%	9%	8%	11%	10%	9%
Below R2	8%	8%	12%	6%	8%	6%	6%	5%	11%	7%
Above R1	8%	14%	12%	12%	11%	14%	15%	14%	11%	13%
Below R1	8%	7%	13%	14%	11%	11%	11%	16%	15%	13%
Above Spot	9%	13%	12%	14%	12%	7%	14%	16%	8%	11%
Below Spot	13%	9%	9%	9%	10%	16%	13%	5%	13%	12%
Above S1	12%	6%	12%	9%	10%	9%	8%	11%	8%	9%
Below S1	17%	20%	7%	14%	14%	13%	10%	7%	13%	11%
Above S2	10%	5%	5%	9%	7%	9%	10%	12%	5%	8%
Below S2	12%	14%	12%	10%	12%	6%	5%	5%	7%	6%
Hypothesis tests on proportions in range:										
Above R2 - Below R2	-1.65	-0.80	-1.53	-0.85	-2.44	0.72	0.85	1.56	-0.31	1.12
Above R1 - Below R1	0.00	1.72	-0.19	-0.71	0.30	0.76	1.12	-0.40	-0.98	0.17
Above Spot - Below Spot	-1.02	0.83	0.61	1.50	1.00	-2.17	0.47	2.53	-1.53	-0.62
Above S1 - Below S1	-1.08	-3.28	1.29	-1.33	-2.32	-0.80	-0.74	0.99	-1.69	-1.41
Above S2 - Below S2	-0.61	-2.24	-1.78	-0.41	-2.46	0.93	1.63	1.77	-0.60	1.75
% out of range										
Higher	11%	16%	16%	11%	14%	18%	15%	26%	14%	17%
Lower	17%	23%	24%	15%	20%	10%	18%	17%	12%	14%
Total	28%	39%	40%	26%	34%	28%	33%	43%	25%	31%

Aggregating across all currencies, there is evidence of selling on breakouts below S1 and S2. The only other significant results are that there are (slightly) more sells below R2 than above, and slightly more buys above S2 than below. This might be evidence of reversal trading, on the belief that prices will move back into the (S2, R2) range. However, the proportion of trades involved is small (5%-8%). Our conclusion is that there is some evidence of simple support/ resistance based trading for Sterling, but not for the other currencies. In general, if the reported support and resistance levels are indeed being used by traders, they are not being used in a simple way.

One feature of trader behaviour that is clear consistent across all recommended trades is the application of target and stop-loss levels. The role of active position management of this kind is well reviewed and illustrated in Lyons (1998) and Bensaid and De Bandt (1998), but has otherwise received little attention from academics. Table 10 shows for all currencies how these limits were used in the management of the trading positions. The great majority of positions that become profitable are closed short of the announced target, or at the target itself. Only 10-15% of winning trades are allowed to run beyond the target price. Conversely, most losing positions are closed at (but rarely before) the stop loss level. About 30% of losing Sterling and Yen trades are closed beyond the stop level, but only 10% of losing Euro and Swiss Franc trades. This conforms well to the traders mantra "cut losses and take profits".

The stop loss limit is much closer to the entry price than the target. The average target is ranges from 83 to 122 ticks from the spot, depending on the currency. The average stop is only 35-45 ticks away. The tight stop-loss limit is another reason why so many of the recommended trades are closed within the hour. Disciplined trading on these limits means that on average winning trades make about twice as much profit as losing trades. Across all currencies, there are about the same number of losing trades as winning trades (more in the case of the Swiss Franc, less in the case of Sterling and the Yen). So the net effect of the tight stop-loss policy is to translate an indifferent directional trading performance into a substantial trading profit. If losses had been

allowed to run as far as profits, the recommended trading positions would not have been profitable.

Table 10. Recommended trades: close, stop loss and target levels

This Table shows characteristics of the trades recommended by analysts – how many trades were winning and how many losing, and the average gain and loss on these trades. The Table also shows where trades were closed relative to the target price and stop loss level of price quoted by the analyst at the time the trade was initiated.

Currency:	EUR	GBP	CHF	JPY
Number of positions:				
Winning	167	224	182	262
Losing	168	194	197	198
Even	16	37	15	-5
Ticks gained/ lost				
Average win	57	66	92	48
Average loss	-31	-32	-42	-29
Total profit (ticks)	4235	8603	8416	6797
Compound value of \$1	1.5969	1.7927	1.6391	1.8074
Average difference (ticks) between				
Target - Open	88	95	122	86
Stop - Open	-35	-39	-45	-37
Winning trades closed:				
Short of Target	104	109	79	136
At Target	38	26	55	20
Beyond Target	20	25	21	18
Losing trades closed:				
Short of Stop	1	0	3	2
At Stop	144	119	161	101
Beyond Stop	23	39	13	48

It is tempting to think that the poor performance of the mechanical rules might be rescued by imposing similar asymmetric stop loss and limit orders on the positions they generate. This is not the case. On the contrary, for almost all (rules x currency) combinations, imposition of a (-.25%, +.5%) stop loss and limit order regime resulted in lower than higher profitability. Good position management has enhanced the profitability of the recommended trades. But the superior performance of the recommended trades relative to the mechanical rules reflects better market timing ability, and not simply better position management.

5.7 CONCLUDING COMMENTS

The aim of this chapter has been to look critically at the way academic studies have evaluated technical analysis. Some conjectures in the academic literature are supported by our study. For example, the directional accuracy of analysts' forecasts is poor. Directional accuracy and mean square accuracy are uncorrelated with profit performance across currencies. The presence of many round numbers in published support and resistance levels, and at trading entry and exit points, gives support to the view that technical factors cause the clustering of exchange rates at round numbers.

Overall, however, our findings cast doubt on the value of many empirical studies of technical trading rules. Following the trading recommendations of the group of analysts studied here would have yielded consistent profits, after allowing for slippage and transactions costs. In contrast, none of the synthetic rules we examined yields significant profits. It would be naïve to think that in a speculatively efficient market all technical analysts could be as successful as those in our sample. But we should recognise the danger that a mechanical interpretation of technical indicators may systematically understate their value, relative to what can be achieved by professional analysts combining technical indicators with other elements of judgement.

None of mechanical rules looks “real”. To varying degrees, they fail to capture key characteristics of the analysts recommended trading positions. Filter rules and directional trading strategies involve excessively long holding periods, and the profits across markets from these rules bear no relation to profits achieved by analysts. Support and resistance breaking rules oversimplify the way that these levels are used in practice, and again fail to reflect the level and pattern of potential profits from trading on technical indicators.

The findings of the last section of the chapter suggest some more fruitful avenues of research on technical analysis. The relationship between the analysts’ trading position and support and resistance levels is complex but does have some logic. Sometimes trades are undertaken after support or resistance is broken, with the implication that prices are expected to continue their trend. Sometimes, trades are undertaken when support and resistance levels are *not* broken, with the implication that prices are expected to reverse their local trend. How analysts can distinguish one situation from another, and more generally whether their behaviour is consistent enough to be captured by an expert system, are interesting questions for research. Another issue concerns the role of limit orders. Much attention is given in academic studies to directional forecasting and decisions over position-taking, but little attention is paid to the role of position management, and the determination of stop loss and limit levels. All of our analysts’ trades are circumscribed by tight stop loss and limit orders. It is interesting to ask how they are determined in practice, and they might be ideally determined.

CHAPTER 6

JUDGMENTAL BOOTSTRAPPING OF TRADING STRATEGIES IN THE BOND MARKET

6.1 INTRODUCTION

In 1954 Paul Meehl (1954/ 1996) published an influential study reviewing 20 pieces of research that compared decisions made by human experts with decisions made by simple linear statistical models. The decisions related to fields as disparate as the diagnosis of schizophrenia, the probability of released prisoners reoffending, and the academic attainments of college students. In every case the statistical model performed as well as, and generally better than, the human judges. This in spite of the fact that the information available to the human judges was usually greater than the limited number of quantitative inputs available to the statistical model. Meehl's work stimulated a fierce and negative reaction from medical practitioners and a summary of their arguments is made in Grove and Meehl (1996). Meehl's work also stimulated further studies widening the domain of comparison between expert judgment and formal models. The meta-analysis by Grove et. al. (2000) found 136 studies, including some in finance-related areas such as bankruptcy prediction and credit rating by banks. Of these studies, 64 showed the statistical approach of weighted linear prediction – multiple linear regression - to be superior, 64 showed approximately the same outcome from human and statistical approaches, and only 8 favoured the human judges.

The aim of this paper is investigate whether simple statistical models can also outperform technical analysts in generating profits from trading bond futures. Technical analysts believe that patterns in the time series of prices can be used to identify profitable trading opportunities. The exercise promises to be interesting for four reasons. First, there seems to be no published evidence on the validity of financial market trading systems built on the performance of experts. The many

books, journals and articles on expert systems in finance are instead concerned with finding complex nonlinear models to determine “ideal” trading positions. Second, technical analysts use a combination of quantifiable indicators, more subjective pattern recognition tools, and the flow of economic and business news to support their decisions. Only the technical indicators can be readily used as inputs to a statistical model, so in this field the human judges have a substantial information advantage. Third, the criterion of success in the markets is profit rather than the percentage of correct outcomes, the typical metric used in the studies cited above. Because of the non-normal distribution of price movements in financial markets, profitability and accuracy are only weakly related, as demonstrated in the study of interest rate forecasts and futures trading by Leitch and Tanner (1991). Fourth, in contrast to, say, corporate bankruptcy events and loan defaults, price movements in financial markets do not have clear-cut drivers. On the contrary, the prediction of the mainstream modern theory of finance is that price changes are made random and unpredictable by the actions of profit-motivated traders.

Indeed, one barrier to the development of expert financial market trading systems is the paucity of objectively verifiable experts (as opposed to self-proclaimed experts, of whom there are many). Section 2 of the chapter introduces the German bond futures markets and two analysts who followed these markets in the years 2000-1, and reports results of tests showing that these analysts do have genuine expertise.

In Section 3 we report the results of a survey designed to establish what technical indicators these (and other) analysts use. Many of the more popular indicators are highly subjective or hard to quantify but some, such as “moving averages” and “stochastics”, can be input to a statistical model. However, they do need to be more precisely defined – e.g. a 20-period moving average, a 5-period stochastic. To control for data-peeking biases, we optimise the parameters of the technical indicators over a pre-test data set consisting of the first 1000 hours of a total data set of 5587 hours during which the relevant market – the Eurex futures exchange - was open.

We then relate the recommended trading positions (short, out, long) to these indicators using data from the next 3000 hours, using the ordered response model of Aitchison and Silvey (1957). In the clinical literature this procedure of modelling experts by simple quasi-linear models is termed the “statistical” or “actuarial” approach. Both terms seem too vague, and we have followed current practice in calling our procedure “judgmental bootstrapping”, the terminology of Dawes (1971), Dawes and Corrigan (1974) and Armstrong (2001). Even then, there is potential for confusion with statistical bootstrap inference methods, compounded here by the fact that we use a methodology related to the statistical bootstrap when testing for analyst expertise.

The results are reported in Section 4 of the paper. Consistent with the findings of the clinical literature, we find that trades based on models of analyst recommendations are more profitable than trades based on the recommendations themselves. Among the alternative specifications investigated, we find that the models that in statistical terms perform best in-sample do yield the highest profits out of sample. However, the pattern of model trades is very different from that of analyst trades, with the models trading more often, and holding positions for much longer. This introduces additional volatility and liquidity risk, and on a risk-adjusted basis it is not at all clear that the models represent an improvement over human judgment.

6.2 DATA AND EXPERTS


Our focus is on the trading recommendations for the Schatz, Bobl and Bund futures contracts in the years 2000-1, made by two analysts then working at Standard and Poor’s MMS (now MMS International). We refer to the analysts by the acronyms FAB and SKY. This service provides continuous commentary and analysis on financial markets, and is distributed worldwide to money market and foreign exchange dealing rooms via the screens of all major quote vendors, including Bloomberg, Reuters and Telerate, and through the company’s website www.globalmarkets.com.

The Schatz, Bobl and Bund are short (approximately 2 year), medium (5 year) and long dated (10 year) German government bonds. There are active markets in futures on these instruments, especially the Bund. Trades are conducted mainly through the German-based EUREX exchange, though the contracts are also traded in other international markets such as the Chicago Mercantile Exchange. Buying or selling a contract effectively fixes a price at which the cash market instrument can be bought or sold at the expiry date of the contract. Like most exchange traded financial futures contracts, Schatz, Bobl and Bund futures expire in March, June, September and December each year. The MMS analyst trading recommendations relate to the nearby contract, and switch to the next delivery month when the nearby month is entered.

Bond futures prices are inversely related to current and expected money market interest rates, with the prices of longer dated bonds reacting to yield changes more sharply than those of shorter dated bonds. Figure 1 shows the time series of nearby Schatz, Bobl and Bund futures prices, aggregated to an hourly frequency from tick-level (transactions) data supplied by *Tickdata.com*. The rising trend in prices reflects the progressive fall in Euro interest rates over the period, as the economic slowdown in the Euro area prompted an easing of monetary policy by the European Central Bank.

The futures contracts themselves are for 1000 bonds each with redemption value €100, and prices of the bonds are quoted to two decimal places as, for example, €105.67. A one tick move in price, from €105.67 to €105.68 say, represents a change of 1 basis point ($1/100^{\text{th}}$ of 1 per cent). Gains and losses on futures contracts are usually measured in ticks or basis points. For the holder of a long position in one contract, a 1 basis point (bp) rise in price would represent a profit of $1/100^{\text{th}} \times 1\% \times €100 \times 1000 = €10$.

Figure 1. Schatz, Bobl and Bund Futures prices, 2000-1.
Source: Eurex and Tickdata.com



Source: Eurex and Tickdata.com

Shortly after the opening of each trading day in London the MMS analyst writes a comment on technical features of the market for each major currency. The idea is that this will inform trader behaviour as the day progresses. Then as prices move through the day and trading opportunities present themselves, the analyst may publish trading recommendations. For example, here are the recommendations and outcomes for the Bund made by SKY who was following the market on 26 May 2000:

Time	Price (end hour)	Analyst	Position Open	Entry	Target	Stop-loss	Close	Profit (bps)
11:00	105.15	SKY	open long	105.11	105.42	104.95		
12:00	105.08							
13:00	104.98	SKY	close long				104.95	-16
14:00	105.02	SKY	open long	105.00	105.32	104.84		
15:00	105.07							
16:00	105.25							
17:00	105.31	SKY	close long				105.32	32

Some time between 10:00 and 11:00 hours London time the analyst suggests buying the Bund contract at a screen price of 105.11. At the same time as making this recommendation, the analyst also gives a “target” price (105.42) that he expects to achieve, and a stop-loss level (104.95) at which he recommends closing the position if the price falls rather than rises. In the event, the price does fall and before 13:00 hours he recommends closing the losing long position at the stop-loss level of 104.95. This trade has lost 16 basis points ($104.95 - 105.11$). Within the next hour the analyst again recommends buying the Bund, at the screen price of 105.00, with a target of 105.32 and a stop loss level of 104.84. This time things go well, and before 17:00 the price has risen to the target level, at which point the analyst recommends closing the position. This trade has made a profit of 32 bps ($105.32 - 105.00$), giving an overall profit of $32 - 16 = 16$ bps for the day.

These trades are representative of the whole data set in the sense that the gap between the entry price and the target is about twice the gap between the entry price and the stop loss level. Analysts are also very disciplined about closing losing positions when stop loss levels are hit. On the other hand, winning positions are sometimes closed short of or beyond the initial price target.

The raw material for our judgmental bootstrapping model consist of data like this through all Eurex trading days in the years 2000 and 2001. Table 1 shows the number of trades recommended by each analyst, the number of winning and losing trades, the profits and losses made on these trades, and the standard deviation of profits across all trades. Over 200 trading recommendations were made for each market. Analyst FAB is the lead analyst for the Schatz and Bobl markets, while SKY follows the Bund. They provide cover for each other at vacation times, so around 10% of all recommendations are not made by the lead analyst. There is some difference in performance across markets. About 60% of recommendations in the Schatz and Bobl are profitable, and about 35% loss-making, as against 51% profitable and 46% loss-making for the Bund. On the other hand, within each market there is little difference

between the two analysts, so the apparent gap in performance between FAB and SKY may reflect the greater efficiency of the more liquid Bund market, and its greater price volatility.

Table 1. Analyst Performance, 2000-1

Market/ Analyst	Recommendations				Average			Standard Deviation	Total Profit
	Total #	Gain #	Even #	Loss #	Gain bps	Loss bps	Profit bps		
<i>Schatz</i>									
FAB	186	113	12	61	9	-7	3	8	571
SKY	16	10	0	6	11	-6	5	9	75
Both	202	123	12	67	9	-7	3	8	646
<i>Bobl</i>									
FAB	228	135	10	83	17	-13	5	16	1220
SKY	23	14	2	7	16	-11	6	14	142
Both	251	149	12	90	17	-13	5	16	1362
<i>Bund</i>									
FAB	34	20	2	12	21	-16	7	19	222
SKY	223	110	6	107	25	-16	4	23	995
Both	257	130	8	119	24	-16	5	22	1217

Notes. For each market and analyst, the table shows the total number of recommended trades, and the number of trades that made gains, broke even, or made a loss. Average gains and losses, the standard deviation of profits across all trades, and total net profits on these trades are measured in basis points (bps). Data is from all trading days in the years 2000-1, and profits are the difference between entry and exit prices reported by the analysts.

We can calculate monetary gains and losses on these trades if some assumption is made about the size of positions. We assume that a position in 1 contract is taken in response to all recommendations. For both analysts the average gain on the profitable trades exceeds the average loss on loss-making trades, so with equal positions on all trades, both analysts generate profits in all markets. Highest profits were achieved in the Bobl, followed by the Bund and the Schatz. Average profit per trade is virtually identical for the two analysts. FAB made 448 trades to gain 2013 bps, an average of

4.5bps per trade. SKY made 262 trades to gain 121 bps, an average gain of 4.6 bps per trade.

Before we proceed to modelling their trading recommendations it is worth asking two questions. First, are the reported profits genuine and robust? Second, is there evidence that the analysts display expertise beyond what could be expected from some naïve trading strategy?

6.2.1 THE ROBUSTNESS OF PROFITS

The question of whether the profits from analyst recommendations are robust is important. The “trades” reported are recommendations rather than transactions conducted in real time, and we need to establish whether it would have been feasible to open and close positions in the way suggested. Because we have tick data, it is straightforward to check whether market transactions occurred at the prices reported by the analysts, and we can pin down the precise times at which these prices were hit. In all cases the prices cited by the analysts were achieved in the market around the time of their published trading recommendations. However, it is unrealistic to imagine that the analyst, or the analyst’s clients, could trade at exactly these prices. One issue is transactions costs, though in futures markets these costs are low. A round trip cost of about €2 per contract (0.2 bps) would be typical for a professional trader. This is low relative to the claimed profits of over 5 bps per trade.

A more important issue is “slippage” – price movements that occur after the trading opportunity has been recognised but before a position can be opened. To measure this we have recalculated profits from recommended trades, but not using the entry and exit prices quoted by the analyst. Instead we imagine trading at the price observed at the end of the 5-minute interval that the recommendation was made, and at the end of the 1-hour interval that the recommendation was made. The results are shown in Table 2. We assume that 1 contract is traded following each recommendation, and

ignore transactions costs. Even a small delay reduces profits substantially. For example, if trades take place at the closing price of the 5-minute window during which the price quoted with recommendation was made (implying an average delay in dealing of $5/2 = 2\frac{1}{2}$ minutes), then profits fall by 20-35%. If we assume dealing at the end of a 15 minute window profits fall further, and if end-hour prices are used profits are only about 60% of those claimed in all markets. This does suggest that the some of the entry and exit prices represent rather transitory trading opportunities, and that analysts' reported profits exaggerate what could be achieved in practice by following their recommendations.

Table 2. Effect of slippage on profitability of recommended trades


	<i>Dealing window</i>			
	<i>None</i>	<i>5 min</i>	<i>15 min</i>	<i>1 hour</i>
Schatz (bps)	646	537	466	390
% of reported profit	100.0%	83.1%	72.2%	60.4%
Bobl (bps)	1362	942	906	795
% of reported profit	100.0%	69.2%	66.5%	58.4%
Bund (bps)	1217	801	863	743
% of reported profit	100.0%	65.8%	70.9%	61.1%

Notes: The table shows profits in basis points over the full sample from executing the analyst recommended trades at the prices reported by the analyst, and at the end of the 5 minute, 15 minute and 1 hour window in which the recommendation was made.

Nonetheless, even allowing for considerable slippage, the recommended trading positions remain comfortably in profit. To simplify our model-building, we assume that both the analysts and the statistical models of analysts trade at end-hour prices. All of our subsequent work is based on this assumption. In the case of the two trades on 26 May, this would have resulted in a loss of $100 \times (104.98 - 105.15) = 17$ bps on the first position, and a gain of $100 \times (105.31 - 104.02) = 29$ bps, making a net profit of 12bps for the day rather than the claimed 16 bps.

Figure 2 shows the path of profits calculated on this basis for the Schatz, Bobl and Bund. In the case of the markets followed by FAB, the Schatz and the Bobl, profits accumulate slowly and steadily across the period. To maintain futures positions, traders must post margin on a daily basis to cover any losses so drawdown, defined as the worst run of losses during the whole trading period, is a key measure of risk. In the case of the Schatz, there is little drawdown. Bobl profits are a little more volatile, with a maximum drawdown of around 100bps in the period January-February 2001. However, the profits generated by SKY in the Bund market are far from even, and show some periods of substantial liquidity risk. Bund trades are highly profitable in the first 6 months of our sample, and again in the last two months, but between these points profits are flat. In this flat period, the Bund trading position is highly volatile, with a drawdown of over 200bps in the months June-August 2000, and over 150bps between December 2000 and January 2001. The Bund profile is especially challenging for our study, since the judgmental bootstrap models are parameterised on data from the middle of our sample period – the shaded area in the Figure – a time when SKY was generating negligible profits.

Figure 2. Analyst profits from futures trades, 2000-1.



Source: Calculated from analyst recommendations, assuming 1 contract traded at end-hour prices.

6.2.2 TESTS FOR EXPERTISE

We have checked the expertise of our analysts in two ways. First, we have looked at the directional accuracy implied by their recommendations – did prices rise after a long position was recommended more often than would be expected by chance? Second, we have looked at the level of profits achieved by our analysts. Do they lie above the range of profits that might be achieved by traders with no skill?

In Table 3 we have set out contingency tables showing for each market whether, when the analysts recommended short or long positions, the price subsequently sent up or went down during the period the position was held. Cases when the market remained unchanged have been ignored, and we have pooled the recommendations of the lead and back-up analyst.

The directional accuracy of the forecasts is not encouraging. The percentage of correct calls is only just above 50% for all markets. We have also compared the actual trades with no-skill benchmark patterns in which the number of long and short positions taken are the same, but there is no correlation between the recommended position and the market direction. For example, in the Bund market the expected number of cases when a short position coincided with a market fall would be $(127 \times 119)/249 = 63.7$, almost exactly the observed number. Not surprisingly, the chi-squared test reflects that the hypothesis is that the distribution is random and it is accepted.

However, poor directional accuracy need not be associated with poor profitability. What matters is “market timing”. Price changes in futures markets are not normally distributed – large moves up and down occur more often than would be expected – and what determines profits is whether the analysts succeed in anticipating these big moves.

Table 3. Directional accuracy of recommended trades

Schatz

Recommended:	Actual		Total
	Down	Up	
Short	32	55	87
Long	35	68	103
Total	67	123	190

Directional Accuracy = 0.52632
Chi-squared(1) = 0.16
p-value = 0.69

Bobl

Recommended:	Actual		Total
	Down	Up	
Short	43	71	114
Long	49	79	128
Total	92	150	242

Directional Accuracy = 0.47107
Chi-squared(1) = 0.01
p-value = 0.93

Bund

Recommended:	Actual		Total
	Down	Up	
Short	61	66	127
Long	58	64	122
Total	119	130	249

Directional Accuracy = 0.45382
Chi-squared(1) = 0.01
p-value = 0.94

Notes: The contingency tables show the number of times the market rose/ fell over the period the analysts recommended long/ short positions in each market. No-change cases have been ignored. Direction accuracy measures the percentage of correct (profitable) trades. The Chi-squared(1) statistic tests for significant differences between the actual pattern of trades and a "random" benchmark which contains the same number of long and short recommendations, but these are uncorrelated with the market going up or down.

One simple test is to compare the profits from the recommended trades with the profits from a naïve “hold and roll” strategy. We have seen that bond futures prices rose through 2000 and 2001. Are the analysts profits greater than would be achieved by buying (say) a March 2000 futures contract at the start of the period, then rolling this into the June, September etc. contracts, eventually selling the December 2001 contract at expiry? These hold-and-roll positions give profits of 144, 357 and 310 bps in the Schatz, Bobl and Bund respectively. Table 2 shows the analysts made 646, 1362 and 1217 bps, and following their recommendations would have made 390, 795 and 743 bps after slippage to the end of each hour, all comfortably higher than the naïve trading rule. The hold and roll strategy is also substantially riskier, since it is continuously in the market and holds positions overnight. This would lead to a drain on cash to meet margin calls during those periods when prices are falling. In contrast, analysts have open positions for only about 30% of all trading hours. Of these, 45% of positions are closed within the same day, over 70% are closed by the end of the following day, over 90% within 3 days. The maximum holding period is 5 days.

To test more formally for market timing ability we conduct a randomisation test on each market. This involves shuffling the trading positions recommended by the analysts 1000f times, and mapping the distribution of profits from the shuffled trades. If the profits achieved by the analysts lie in towards the upper tail of this distribution (say at the p -th percentile), we can reject the null hypothesis that the analysts performance is equal to that of a random trader at the $(1-p)\%$ level. Randomisation tests (sampling without replacement from the set of trading positions) are more appropriate than the popular bootstrap procedure (sampling with replacement) in testing hypotheses about relationships between variables. The bootstrap is directed at refining estimates of population parameters. Our shuffling procedure ensures that in each replication the “random trader” has exactly the same balance of long and short positions, and exactly the same distribution of hours spent in each position, as the analyst.

Table 4 reports the percentiles of the profits achieved by the analysts separately over three sub-periods of our data – the “pre-test” period (hours 1-1000) on which our technical indicators will be developed; the “in-sample” period (hours 1001-4000) on which our bootstrap models are parameterised; and the “out-of-sample” period (hours 4001-5587) on which forecasting performance will be evaluated. In the pre-test period the analyst profits exceed 99th percentile of profits from randomised trades. In the in-sample period the Schatz and Bobl trades continue to perform at this level. But the weak performance of the Bund trades in mid-sample noted in Figure 2 means that they only exceed the 80th percentile of the randomised distribution. In the out-of-sample period the Schatz and Bobl exceed the 95th percentile, and the Bund exceeds the 90th percentile of the randomised distribution. On this basis we can conclude that the Schatz and Bobl show expertise at conventional levels of significance. Evidence on the Bund is less consistent over time, but overall suggests better than random market timing ability.

Table 4. Bootstrap analysis of profits from recommended trades

<i>Sample</i>		<i>Schatz</i>	<i>Bobl</i>	<i>Bund</i>
<i>Pre-test</i>	Profit (bps)	101	337	291
	Percentile	100	100	99.2
<i>In-sample</i>	Profit (bps)	221	220	14
	Percentile	99.8	100	84.6
<i>Out-of Sample</i>	Profit (bps)	68	238	438
	Percentile	90.6	94.4	94.6
<i>Whole period</i>	Profit (bps)	390	795	743
	Percentile	96.2	98.7	91.4

Notes: the table shows profits in basis points from analyst trades in the pre-test period (hours 1-1000), the in-sample period (hours 1001-4000) and the out-of-sample period (hours 4001-5587). These profits have been compared with the distribution of profits from 1000 random shuffles of the analyst trading positions.

6.3 TECHNICAL INDICATORS

While all technical analysts agree that market psychology is revealed in the pattern of prices, there are many different approaches to uncovering these patterns. As discussed in our literature survey in Chapter 3, some analysts use well defined indicators based on averages or ratios of recent prices. Some use recent data to identify normal ranges for prices and trade on breakouts above the upper “resistance” level or below the lower “support” level. Some rely on less well defined patterns – such as “double tops” and “head-and-shoulder formations” – that are claimed to anticipate price reversals. Others rely on yet more vague “wave theory”, “Gann numbers” and “Fibonacci ratios” to determine market trends and likely turning points. Almost all analysts follow more than one indicator, and “confirmation” of price changes is emphasised in popular technical analysis texts such as Edwards and Magee (2001), Murphy (1998) and Pring (1998).

We are not concerned here with the merits of these indicators, though it is worth noting that evidence from academic research is mixed. Some positive results applying statistical indicators and trading channels to the stock market appear in Brock, Lakonishok and Schleifer (1992), though their conclusions are disputed on data-mining grounds by Sullivan, Timmerman and White (1999). Levich and Thomas (1993) report profits from support/ resistance type rules (filter rules) in the foreign exchange market, though again the persistence of these over time has been questioned (Olson, 2003). Pattern recognition methods have been less studied, possibly because they are less well defined. Lo, Mamaysky and Wang (2000) find that typical technical analysis patterns occur more often than chance, and do anticipate changes in the statistical properties of the price series. However, they stop short of claiming that the patterns can be used to generate profits, and Chang and Osler (1999) find no profits from trading on the “head-and-shoulders” pattern. As tools for understanding how technical analysts behave, these studies are not especially insightful. All the indicators are generated *ex post* by researchers rather than *ex ante* by technical analysts, and all look at single indicators in isolation.

6.3.1 SURVEY OF METHODS USED BY ANALYSTS

We do, however, have some insight into how the MMS analysts work. Through 1998 we conducted a weekly survey of foreign exchange and stock market forecasts made by 14 MMS analysts, including FAB and SKY. Full results are reported in Batchelor and Kwan (2001) and chapter 4 in this thesis. As part of this exercise we asked each week and for each target variable what “percentage weight” the analyst had given to different technical indicators in framing their forecast. All analysts reported using more than one indicator. Different analysts did favour different methods, but for individual analysts there was almost complete consistency over time in the ranking of the methods they used. Table 5 shows the average (over time) weights assigned by our two analysts FAB and SKY, and the average (over time and analysts) responses of the other 12 respondents.

Table 5. Relative popularity of technical indicators among MMS forecasters

Technique:	Forecaster:		
	FAB	SKY	12 Others
Support and Resistance	8.2%	2.7%	19.4%
Chart Patterns	5.9%	30.0%	19.0%
Stochastics	37.5%	21.8%	14.1%
Fibonacci Numbers	13.1%	1.9%	13.2%
Moving Averages	0.4%	1.6%	12.8%
Channels	11.2%	1.6%	6.7%
Relative Strength	0.8%	1.9%	4.5%
Others	1.4%	13.6%	3.9%
Elliott Wave	0.0%	22.2%	3.8%
Candlestick Charts	21.6%	2.7%	2.7%

Notes: Based on weekly surveys among MMS forecasters in the period 9 January – 6 November 1998. Percentages are the “weights” that analysts assigned to the different technical indicators in framing their forecasts, averaged over the 44 weeks of the survey.

The most popular techniques among the whole group are pattern recognition methods, based on chart patterns and support and resistance “trendlines”. Some statistical indicators – stochastics, moving averages, and to a lesser extent the relative strength index - are also popular, though the use of Fibonacci ratios to fix price targets and reversal levels was given more weight than moving average methods. The two analysts studied here do not conform exactly to the typical pattern. Both place some importance on stochastics, and FAB uses trendlines and channels quite heavily. But each has an idiosyncratic style. FAB is a devotee of Japanese candlestick charts, for example, which involves recognising patterns in a sequence of box-plots of open/ high/ low/ close data. Nison (1991) popularised these methods in the West in the 1990s. On the other hand SKY is a devotee of the Elliott Wave (Prechter and Frost, 1998), which involves identifying where the current market price lies within sets of five rising and falling waves.

On the face of it, this is not encouraging for a bootstrapping exercise, since it suggests that analysts in practice place most weight on hard-to-formalise chart patterns. It is possible to automate chart pattern recognition - see for example the studies by Lo et. al. (2000) and Chang and Osler (1999) cited above - but there are many possible patterns, and identification is a complex process. This is especially true of esoteric methods like candlestick charting and the Elliot Wave. However, the idea of judgmental bootstrapping is to mimic analysts decisions using a small number of easily quantifiable inputs. We are not trying to mimic how they reach these decisions. So while priorities assigned by the analysts to different indicators are suggestive (we will certainly employ a stochastic, for example), they need not limit the choice of inputs to the bootstrapping model.

We use five conventional technical indicators designed to predict turning points in the price series, and a set of recent open/ high/ low close prices intended to summarise the information in a candlestick chart. All either were or could have been available to the analysts at the time they made their trading recommendations.

6.3.2 TECHNICAL INDICATORS

The conventional technical indicators are:

PMA : the ratio of the closing price to a moving average of past closing prices, defined as

$$PMA_t = P_t / \sum_{i=0}^{k-1} P_{t-i}$$

where P_t is the price at time (hour) t . The idea is that as the price cuts the moving average from below, this signals reversal of trend from falling to rising, and vice versa. So positive values of $X1$ should be associated with long positions, and negative values with short positions. The optimum bandwidth k for the moving average remains to be determined.

2. SLMA: the ratio of a short ($k1$ period) moving average to a longer period ($k2$) moving average, defined as:

$$SLMA_t = \sum_{i=0}^{k1-1} P_{t-i} / \sum_{i=0}^{k2-1} P_{t-i}$$

The rationale for this is similar to the simple moving average rule. The simple rule is subject to “whipsawing”, frequent reversal of position when the market is not trending. The short/ long moving average rule is less vulnerable to whipsawing, but achieves this at the cost of delayed signals when the market really has changed direction. In Figure 3 we show a short term moving average (SMA) and a long term moving average (LMA) overlaid on the price series. Buying and selling might be triggered at crossovers of the two series.

3. MOM: the k -period “momentum” in the market, defined as

$$MOM_t = 100 \times P_t / P_{t-k}$$

If the price decelerates as it reaches a peak, this momentum measure will fall from a high positive value (above 100), and the fall through some critical “overbought” level can be used as a signal to close a long position. Conversely if the fall in the market loses momentum near the trough, an increase in MOM from low values through some critical “oversold” level can be taken as a signal to close a short position.

4. RSI: the “Relative Strength Index” defined as

$$RSI_t = 100 - 100 / (1 + MAUP_t / MADOWN_t)$$

where $MAUP_t$ is the average of the price changes in the last k periods in which prices closed up (i.e. for which $P_{t-i} > P_{t-i-1}$), and $MADOWN_t$ the average of the price changes when the market closed down. As the market approaches a peak, the up-closes tend to be less marked, and the index will start to fall from a high value. As it falls through some overbought level, this would be taken as a signal to close a long position.

5. STOCH: the “Stochastic” indicator, popular in our survey of analysts, and defined as

$$\%K_t = \frac{P_t - \min(L_{t-k+1}, L_{t-k+2}, \dots, L_t)}{\max(H_{t-k+1}, H_{t-k+2}, \dots, H_t) - \min(L_{t-k+1}, L_{t-k+2}, \dots, L_t)}$$

$$STOCH_t = \sum_{i=0}^{k/2} \%K_{t-i} / k/2$$

where L_t and H_t are respectively the low and high prices observed in the hour between $t-1$ and t . This is used in the same way as the momentum and relative strength indicators. All three try to measure the rate of change of the underlying cycle in prices, and are often termed “oscillators”. The rationale for the stochastic oscillator is that as the market nears a peak (say), the closing price falls ways from recent highs,

and this can be used as a signal to close a long position. Note that it uses more information than the other indicators, in the form of hourly high and low prices. A stochastic oscillator is shown on Figure 3, along with conventional 80/20 overbought and oversold levels. Buying and selling would be triggered by crossovers of the stochastic with these levels. The fact that the second peak in the stochastic looks weaker than the second peak in the “double top” in the price series might also trigger a sell.

For all indicators we need to make decisions about the number of time periods over which we compute the averages and rates of change required. To eliminate the possibility of data mining, we have selected the lag lengths that would have produced the maximum profits for each indicator taken separately, using data from the first 1000 hours of our data base of 5587 hours. These optimised indicators are then used as inputs to bootstrap models that are parameterised using data from next 3000 observations 1001-4000. The remaining 1587 observations 4001-5587 are used to evaluate the out-of-sample predictions of the models.

The ranges of lag over which we searched in optimising each indicator, and the outcomes of our experiments, are reported in Table 6. For the PMA rule applied the Bund, for example we found maximum profit with rule based on the ratio of the current price to the average of 10 current and lagged prices. Use of this rule (long when $PMA > 1$, short when $PMA < 1$) yielded a profit of 266 basis points, involving 52 trading positions. Similarly, the best Bund SLMA rule is based on the relation between a 4-hour and a 30-hour moving average. And so on. In the case of the oscillators MOM, RSI and STOCH, the critical overbought/ oversold levels were set one standard deviation above and below the mean values of the indicators. The Table shows that for all markets and indicators, there were some lag length(s) that produced positive profits in hours 1-1000. This is no guarantee, of course, that the same lags will be effective out-of-sample.

Table 6 Optimum values for technical indicators

		<i>Parameters (hrs)</i>		<i>Pre-test Profits (bps)</i>
		<i>k1</i>	<i>k2</i>	
PMA	Schatz	6		24
	Bobl	8		90
	Bund	10		266
SLMA	Schatz	5	18	-9
	Bobl	3	36	-39
	Bund	4	30	158
MOM	Schatz	9		59
	Bobl	22		114
	Bund	17		368
RSI	Schatz	9		75
	Bobl	10		198
	Bund	10		661
STOCH	Schatz	7	5	70
	Bobl	9	2	137
	Bund	6	6	381

Notes: The indicators are:

PMA = ratio of price to a k1 hour moving average

SLMA = ratio of a k1 hour moving average to a k2 hour moving average

MOM = momentum, the ratio of the current price to the price k1 hours previously

RSI = relative strength index, based on a window of k1 hours

STOCH = stochastic oscillator, a k2 hour moving average of a k1 hour simple stochastic.

Optimum values for k1 and k2 are those that maximised profits over hours 1-1000 of our sample, in the range 1-40hrs. Profits made by the analysts in this period were 101bps, 337bps and 291bps for the Schatz, Bobl and Bund respectively.

6.3.3 CANDLESTICK CHARTS

Figure 3 also shows a series of hourly candlesticks for the Bund. The range between the open and close price of each hour is represented by a box (the “body” of the candlestick), coloured white if the close was above the open, and black if the close was below the open. Lines are then drawn from the box to the hourly high and low prices, the upper and lower “shadows” of the candlesticks. Two patterns are labeled on the Figure. The “doji” is a candlestick with no body, since the open price and the close price are the same. If a doji is observed after a sustained up- or down-trend, this is taken as a sign that a turning point may occur. The “dark cloud cover” pattern occurs when there is a rising trend in the price, but after rising through one period (a white candlestick) the market then opens higher, but closes down (a black candlestick), the closing price being well within the body of the previous period. This again is taken as a signal that the trend is reversing. Around 70 patterns of this kind, involving sequences of 1, 2 or 3 candlesticks, are described in Nison (1991).

To summarise the information on the size and colour of the body of the candlestick, and the size of the upper and lower shadows, we use the statistics

$$\text{CLOSE}_t - \text{OPEN}_t$$

$$\text{HIGH}_t - \text{CLOSE}_t$$

$$\text{CLOSE}_t - \text{LOW}_t$$

and to capture the information in two-period candlesticks we use lagged values of these figures, together with the change in closing prices

$$\text{CLOSE}_t - \text{CLOSE}_{t-1}$$

Note that candlestick charts are really designed for daily data from markets with well defined opening and closing periods, and not for continuously traded instruments. For our intraday data the close of one period and the open of the next are the same except at the close of each day, and this limits the number of patterns.

6.4 ORDERED RESPONSE MODELS

The target variable for our bootstrap model is the 1-hour ahead position recommended by the MMS analysts, labeled POSIT_{t+1} . This takes the values -1 , 0 and $+1$ according as the recommended position is short, no position, or long. The trivariate distribution of POSIT means that we cannot use a simple ordinary least squares regression model to link it to technical indicators. Instead we apply the ordered-response model of Aitchison and Silvey (1957). This imagines a continuous normally distributed latent variable Y_{t+1}^* which is linearly related to the vector of technical indicators as

$$Y_{t+1}^* = X_t' \beta + u_t \quad (1)$$

where X_t is the vector of technical indicators and candlestick statistics, and β is a vector of coefficients. The discussion of the indicators earlier suggested that for some it is the level of the indicator that matters, and for others the change in the indicator. To admit both these possibilities we define X_t as the vector of indicators observed at t and $t-1$. Since we want to explore the possibility of combining rules from all three markets, all indicators are normalised to have mean = 0 and standard deviation = 1.

The fitted positions are then assumed to depend on the value of the latent variable relative to thresholds γ_{SHORT} and γ_{LONG} as

$$\text{POSITF}_{t+1} = -1 \text{ if } \hat{Y}_{t+1}^* < \gamma_{\text{SHORT}}$$

$$\text{POSITF}_{t+1} = 0 \text{ if } \gamma_{\text{SHORT}} < \hat{Y}_{t+1}^* < \gamma_{\text{LONG}}$$

$$\text{POSITF}_{t+1} = +1 \text{ if } \hat{Y}_{t+1}^* > \gamma_{\text{LONG}}$$

These thresholds are estimated simultaneously with the equation parameters β using maximum likelihood methods.

Tables 7A – 7C shows the results of estimating the ordered response model (1) for the Schatz, Bobl and Bund respectively. The first set of coefficients in the table is for an “unrestricted” model that includes all the possible current and lagged inputs. For the Schatz and Bobl, the indicators include the candlestick open/ high/ low/ close data, since FAB claims to use these methods. The second set of coefficients is for the “restricted” model that results from progressively eliminating statistically insignificant variables from the unrestricted model. This results in a more parsimonious model with current and lagged price/ moving average and short/long moving average indicators, and current momentum and stochastic indicators. The third model includes only current (hour t) technical indicators, and no lagged indicators. The final column contains current variables, but in the case of the Schatz and Bobl omits the candlestick proxies, and in the case of the Bund adds these variables.

Table 7A. Ordered Response Models of recommended trading positions: Schatz

	<i>Unrestricted</i>		<i>Restricted</i>		<i>Current Variables</i>		<i>Current, no Candlesticks</i>	
	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>
<i>PMA(t)</i>	0.2074	0.82	0.3028	4.58	0.0830	1.81	0.0937	2.50
<i>SLMA(t)</i>	-0.3857	-1.21	-0.4657	-1.83	0.2275	5.22	0.1223	2.84
<i>MOM(t)</i>	-0.0106	-0.13			-0.0648	-1.02	-0.0669	-1.06
<i>RSI(t)</i>	-0.0169	-0.31			0.0205	0.70	0.0739	2.54
<i>STOCH(t)</i>	-0.1925	-1.77	-0.1894	-1.82	-0.0482	-1.31	0.0148	0.41
<i>PMA (t-1)</i>	0.0791	0.37						
<i>SLMA (t-1)</i>	0.5179	1.88	0.5667	2.45				
<i>MOM (t-1)</i>	-0.0704	-0.84						
<i>RSI(t-1)</i>	0.0478	0.90						
<i>STOCH(t-1)</i>	0.1800	1.69	0.1722	1.73				
<i>CLOSE-OPEN</i>	-3.5629	-1.62	-3.3743	-2.51	-3.5847	-1.64		
<i>HIGH-CLOSE</i>	0.3129	0.21			0.6953	0.49		
<i>CLOSE-LOW</i>	-0.2873	-0.17			-0.5973	-0.37		
<i>CLOSE-CLOSE</i>	-2.1019	-1.23			3.8815	2.04		
<i>CLOSE-OPEN(-1)</i>	0.2951	0.20	-2.7333	-2.42				
<i>HIGH-CLOSE(-1)</i>	-0.2541	-0.15						
<i>CLOSE-LOW(-1)</i>	2.7157	0.41						
<i>γSHORT</i>	-1.0697	-21.42	-1.0713	-37.50	-1.0662	-24.06	-1.1120	-38.24
<i>γLONG</i>	1.0083	20.34	1.0050	36.08	1.0064	22.92	0.9627	35.21
<i>Statistics</i>								
<i>Pseudo R-squared</i>	0.0191		0.018526		0.0168		0.0142	
<i>AIC</i>	1.6478		1.642172		1.6464		1.6428	

Table 7B. Ordered Response Models of recommended trading positions: Bobl

	<i>Unrestricted</i>		<i>Restricted</i>		<i>Current Variables</i>		<i>Current, no Candlesticks</i>	
	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>	<i>Coefficient</i>	<i>z-statistic</i>
<i>PMA(t)</i>	-0.2773	-0.88			0.0469	0.90	0.0689	1.73
<i>SLMA(t)</i>	0.4052	1.10	0.2905	10.01	0.2043	2.79	0.1898	2.82
<i>MOM(t)</i>	0.1529	1.29			0.0842	1.21	0.0970	1.48
<i>RSI(t)</i>	-0.0459	-0.87	-0.0586	-2.27	-0.0582	-2.06	-0.0533	-1.90
<i>STOCH(t)</i>	-0.0711	-0.94			-0.0167	-0.41	-0.0308	-0.81
<i>PMA (t-1)</i>	0.2492	0.97						
<i>SLMA (t-1)</i>	-0.1877	-0.58						
<i>MOM (t-1)</i>	-0.0825	-0.68						
<i>RSI10(t-1)</i>	-0.0262	-0.51						
<i>STOCH(t-1)</i>	0.0497	0.70						
<i>CLOSE-OPEN</i>	0.4912	0.44			0.4284	0.38		
<i>HIGH-CLOSE</i>	0.2177	0.32			0.0662	0.10		
<i>CLOSE-LOW</i>	-1.1492	-1.51			-1.5947	-2.19		
<i>CLOSE-CLOSE</i>	3.1334	1.03	0.8332	2.13	0.6416	0.62		
<i>CLOSE-OPEN(-1)</i>	1.0706	1.37	1.0642	2.02				
<i>HIGH-CLOSE(-1)</i>	0.1638	0.24						
<i>CLOSE-LOW(-1)</i>	-1.5322	-2.00	-1.8013	-2.68				
<i>γSHORT</i>	-1.1786	-23.19	-1.1562	-27.33	-1.1438	-25.07	-1.0738	-37.39
<i>γLONG</i>	0.9202	18.68	0.9391	23.08	0.9516	21.46	1.0190	36.20
<i>Statistics</i>								
<i>Pseudo R-squared</i>	0.0272		0.0258		0.0258		0.0247	
<i>AIC</i>	1.6344		1.6288		1.6313		1.6306	

Table 7C. Ordered Response Models of recommended trading positions: Bund

	Unrestricted		Restricted		Current Variables		Current, incl Candlesticks	
	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic
<i>PMA(t)</i>	0.0324	0.65	0.0309	0.62	0.0663	1.86	0.0457	1.16
<i>SLMA(t)</i>	-0.6823	-1.74	-0.7914	-2.24	0.6229	9.78	0.6473	9.57
<i>MOM(t)</i>	-0.1752	-1.88	-0.2173	-3.24	-0.2140	-3.35	-0.2377	-3.53
<i>RSI(t)</i>	0.0075	0.13			0.0264	0.88	0.0188	0.62
<i>STOCH(t)</i>	-0.2118	-2.19	-0.1986	-5.31	-0.1517	-4.35	-0.1411	-3.96
<i>PMA (t-1)</i>	0.3184	3.14	0.3466	3.81				
<i>SLMA (t-1)</i>	1.2089	3.54	1.2855	4.09				
<i>MOM (t-1)</i>	-0.0666	-0.68						
<i>RSI(t-1)</i>	0.0023	0.04						
<i>STOCH(t-1)</i>	0.0139	0.16						
<i>CLOSE-OPEN</i>							0.2512	0.29
<i>HIGH-CLOSE</i>							1.0995	1.79
<i>CLOSE-LOW</i>							0.0306	0.05
<i>CLOSE-CLOSE</i>							0.5699	0.79
<i>CLOSE-OPEN(-1)</i>								
<i>HIGH-CLOSE(-1)</i>								
<i>CLOSE-LOW(-1)</i>								
<i>γSHORT</i>	-1.1232	-37.85			-1.1219	-37.93	-1.0521	-20.88
<i>γLONG</i>	1.1507	38.21			1.1420	38.15	1.2148	23.53
<i>Statistics</i>								
<i>Pseudo R-squared</i>	0.0517		0.0515		0.0478		0.0489	
<i>AIC</i>	1.5141		1.5116		1.5169		1.5178	

Notes: Parameter estimates of the ordered response model (1) obtained by applying maximum likelihood to observations (hours) 1001-4000 of our data set, with associated z-statistics, pseudo- R^2 and Akaike information criterion (AIC) values. The dependent variable is the 1-hour ahead recommended trading position $POSIT_{t+1}$, and regressors are current (hour t) and lagged ($t-1$) technical indicators.

Overall, these models do not fit the observed pattern of traders recommendations at all well. The R^2 is only of the order of $1\frac{1}{2}$ - 2% for the Schatz, about $2\frac{1}{2}\%$ for the Bobl, and around 5% for the Bund. Since the Bund is the most volatile price series, this does suggest that the recommendations made by SKY may have more logic to them than the recommendations made by FAB. In all cases, the restricted model is judged by the Akaike Information Criterion to show the best tradeoff of fit against parsimony. The ratio of short to long run moving average (SLMA) features in the restricted models for all markets, and all contain at least one oscillator – the stochastic in the case of the Schatz and Bund, the relative strength index in the case of the Bobl. Our survey found that FAB uses candlesticks but SKY does not. Consistent with this, there are some significant lagged candlestick proxies in the models for the Schatz and Bobl, but not for the Bund model, which instead feature a greater number of conventional indicators.

Given an observed vector of indicators X_T at time T , the probabilities of SHORT, OUT (of the market) and LONG positions at $T+1$ predicted by the ordered response model are

$$P(\text{SHORT}) = \Phi (\gamma_{\text{SHORT}} - X_t' \beta_i)$$

$$P(\text{OUT}) = \Phi (\gamma_{\text{LONG}} - X_t' \beta_i) - \Phi (\gamma_{\text{SHORT}} - X_t' \beta_i)$$

$$P(\text{LONG}) = \Phi (\gamma_{\text{LONG}} - X_t' \beta_i)$$

where Φ is the cumulative normal function. Positive coefficients increase the probability of the high-valued response (LONG position) and reduce the probability of a SHORT position. The estimated coefficients correspond to some extent with the intuitions about the indicators offered above. Looking at the restricted model for the Bund, for example, it is reasonable that there is a positive coefficient on the PMA variable (and a net positive sign on the SLMA variable), since we would expect a long position to be associated with a price that is high relative to some moving average. In the case of the MOM and STOCH oscillators we expect to enter a short position when they are high and falling, and a long position when they are low and rising. This does

rationalise the observed negative coefficients on the levels of the variables, but not the absence of lagged terms.

While it would be comforting to have interpretable coefficients, what matters is whether the statistical model matches or outperforms the profit performance of the traders on which it has been parameterised. The obvious way to transform fitted and forecast values of the models into trading positions is the Maximum Probability Rule – that is, we take the position that has the highest probability according the model. However, this rule in no sense mimics the behaviour of the analysts. Across all the markets, the analysts recommended being in the market in about 30% of hours in the in-sample period about 70% of the time. In contrast, as a direct result of the low R^2 of the models, the Maximum Probability Rule recommends being in the market less than 1% of the time, and so is out of the market almost all of the time.

Since the criterion for success in our experiments is profit, a more natural way of determining positions is the Maximum Profitability Rule, under which we take positions:

SHORT if $P(\text{SHORT}) > \pi_{\text{SHORT}}$

LONG if $P(\text{LONG}) > \pi_{\text{LONG}}$

where $P(\text{SHORT})$ and $P(\text{LONG})$ are as before estimated probabilities from the ordered response model, and π_{SHORT} and π_{LONG} are thresholds. We determine π_{SHORT} and π_{LONG} so as to maximise profits in the in-sample period. The first two rows of Tables 8A – 8C show the results of applying this rule to the Schatz, Bobl and Bund. The probability thresholds are asymmetric, with π_{SHORT} around .20 and π_{LONG} around .10. The resultant distribution of trading positions is again different from that of the analyst, but this time the models recommend more trading rather than less. Indeed, under the preferred restricted model, we are *always* in the market, and have a long position in 80-85% of hours.

Tables 8A-C also show the profits from trading on the basis of the models for in-sample and out-of-sample periods. The restricted model yields profits in excess of analyst profits in all markets and time periods. The excess over analyst profit is smallest in the case of the Schatz (277 v. 221 bps in-sample, 112 v. 68 bps out of sample), the market where the model fit was worst. Model profits are decisively higher for the Bobl and the Bund. and during in the in-sample period for the Bund - when the analyst generated only 14bps - the model trades yield 662bps. Looking across the different models, it seems that out-of sample profits are better predicted by in-sample statistical performance (as measured by the AIC) than by in-sample profits. The statistically preferred restricted model generally performs best or second best out-of-sample. In contrast, the models that make highest profits in-sample – those including only current variables – perform inconsistently out-of-sample.

Table 8A. Trading positions and profits from ordered response model, Schatz

<i>Schatz</i>			<i>Recommended Trades</i>	<i>Model:</i>			
				<i>Unrestricted</i>	<i>Restricted</i>	<i>Current</i>	<i>Current, excl Candlesticks</i>
<i>Maximum Profit Rule</i>							
	π_{SHORT}			0.21	0.17	0.18	0.17
	π_{LONG}			0.10	0.13	0.13	0.13
<i>In-sample</i>	SHORT	(hours)	410	251	646	476	575
	OUT	(hours)	2072	0	0	101	0
	LONG	(hours)	518	2749	2354	2423	2425
	<i>Profits</i>	<i>(bps)</i>	<i>221</i>	<i>301</i>	<i>277</i>	<i>310</i>	<i>305</i>
<i>Out-of-sample</i>	SHORT	(hours)	465	249	537	417	493
	OUT	(hours)	997	0	0	54	0
	LONG	(hours)	125	1411	1222	1247	1253
	<i>Profits</i>	<i>(bps)</i>	<i>68</i>	<i>64</i>	<i>112</i>	<i>17</i>	<i>170</i>

Table 8B. Trading positions and profits from ordered response model, Bobl

<i>Bobl</i>			<i>Recommended</i>	<i>Model:</i>			
			<i>Trades</i>	<i>Unrestricted</i>	<i>Restricted</i>	<i>Current</i>	<i>Current, excl</i>
							<i>Candlesticks</i>
<i>Maximum Profit Rule</i>							
	<i>p</i> SHORT			0.20	0.20	0.20	0.19
	<i>p</i> LONG			0.11	0.11	0.11	0.11
<i>In-sample</i>	SHORT	(hours)	440	461	447	403	437
	OUT	(hours)	2070	83	72	0	0
	LONG	(hours)	490	2456	2481	2597	2563
	<i>Profits</i>	<i>(bps)</i>	220	496	421	495	375
<i>Out-of-sample</i>	SHORT	(hours)	242	572	558	521	537
	OUT	(hours)	1163	27	24	0	0
	LONG	(hours)	182	1228	1240	1283	1276
	<i>Profits</i>	<i>(bps)</i>	238	329	545	527	555

Table 8C. Trading positions and profits from ordered response model, Bund

<i>Bund</i>			<i>Recommended</i>	<i>Model:</i>			
			<i>Trades</i>	<i>Unrestricted</i>	<i>Restricted</i>	<i>Current</i>	<i>Current, incl</i>
							<i>Candlesticks</i>
<i>Maximum Profit Rule</i>							
	π SHORT			0.20	0.19	0.19	0.18
	π LONG			0.08	0.09	0.08	0.09
<i>In-sample</i>	SHORT	(hours)	428	578	622	666	827
	OUT	(hours)	2147	0	57	0	0
	LONG	(hours)	425	2422	2321	2334	2173
	<i>Profits</i>	<i>(bps)</i>	14	597	662	711	571
<i>Out-of-sample</i>	SHORT	(hours)	242	586	629	388	436
	OUT	(hours)	1226	0	24	0	0
	LONG	(hours)	119	1235	1185	1199	1151
	<i>Profits</i>	<i>(bps)</i>	438	551	532	539	455

Notes: The bold figures are profits in basis points from the analysts recommended trades, and from the judgmental bootstrap models. π SHORT and π LONG are the thresholds for the probability of long and short positions that maximise in-sample (hours 1001-4000) profits. The resulting profile of hours long, short and out-of-the market are shown for in- and out-of-sample(hours 4001-5587) periods.

It seems that simple models parameterised on analysts trading recommendations can produce at least as much profit as the analysts themselves, and in most cases substantially more. However, the pattern of model-based trades suggests that this excess profit is not achieved without an increase in risk. Figures 4A-C illustrate the problem. The final level of (restricted) model profits is above that of the analyst in both in-sample and out-of-sample periods in all cases. But because the model is in the market continuously, the volatility of profits is high, and quite large losses can be experienced *en route* to the final higher level of profits.

Figure 4A. In-sample and Out-of-sample profitability of bootstrap model: Schatz




Figure 4B. In-sample and Out-of-sample profitability of bootstrap model: Bobl




Figure 4C. In-sample and Out-of-sample profitability of bootstrap model: Bund




Table 9 brings together the analyst and model profits, and sets them against relevant measures of risk – the number of positions taken, the standard deviation of daily changes in profit, and the maximum drawdown. The problem with model-based trades is not increased frequency of trading – the total number of positions taken is similar to the analysts, so transactions costs would also be similar. The problem is the longer duration of model positions. The daily standard deviation of model-based profits is about twice that of the analyst profits. The maximum drawdown in-sample for the Bund is about the same for analyst and model, but in all other cases drawdown is also substantially higher for model. In practical terms this means that model-based trades would be subject to frequent and large margin calls, and the profit achieved would be sensitive to exactly when trading starts and ends.

Table 9. Return and risk in analyst and model trades

		<i>Schatz</i>		<i>Bobl</i>		<i>Bund</i>	
		<i>Analyst</i>	<i>Model</i>	<i>Analyst</i>	<i>Model</i>	<i>Analyst</i>	<i>Model</i>
Profit <i>(bps)</i>	<i>In-sample</i>	221	277	220	421	14	662
	<i>Out-of-sample</i>	68	112	238	545	438	532
No. of positions	<i>In-sample</i>	84	165	109	137	119	115
	<i>Out-of-sample</i>	50	80	59	62	57	57
Daily SD <i>(bps)</i>	<i>In-sample</i>	1.3	2.7	2.9	5.6	4.2	8.1
	<i>Out-of-sample</i>	2.1	3.2	3.3	7.2	5.1	10.0
Max Drawdown <i>(bps)</i>	<i>In-sample</i>	-26	-99	-133	-170	-237	-248
	<i>Out-of-sample</i>	-45	-108	-62	-250	-75	-408

Notes: The table shows the total profit, number of positions taken, standard deviation of hourly changes in profit, and maximum drawdown for analyst and (restricted) model trades during the in-sample and out-of-sample periods (hours 1001-4000, and 4001-5587). profits and losses are measured in basis points.

6.5 BEYOND THE JUDGMENTAL BOOTSTRAP

In the spirit of the clinical studies reviewed earlier, we have deliberately kept the modelling process simple. It may be argued that this biases our results against finding value in the judgmental bootstrap. So in this final section of the paper we examine whether profit performance can be improved by some extensions of the ordered response model.

First, combining forecasts from different sources has been found to improve mean square accuracy across a range of forecasting tasks (Armstrong, 2001), so a natural question is whether combining will also improve profitability. We have looked at three ways of combining analyst and model information. We combine current analyst and model judgment, and take a position only if the analyst and model agree. We add the most recent analyst decision, the lagged dependent variable $POSIT_t$, as a regressor

in the ordered response model (1). And we pool expertise across the three markets - effectively across both analysts - by using our normalised technical indicators to estimate a single pooled model.

Second, in principle the relationship between technical indicators and trading positions is complicated, and it is possible that the ordered response model is too simple to capture this. To test this proposition, we estimate a single layer neural network of the form

$$\text{POSIT}_{t+1} = Z_t' \delta + u_t \quad (2)$$

where δ is a vector of coefficients ("weights") and Z_t is the vector $[1 \ X_{1t} \ X_{2t} \ \dots X_{nt}]$ of n sigmoid transformations of the inputs. Each X_{it} ("node") has the form

$$X_{it} = \Phi(X_t' \beta_i)$$

where we have taken Φ to be the cumulative normal distribution function, and the β_i are coefficient vectors. A function of this kind can approximate the trivariate distribution of POSIT to any required degree of accuracy, and can account for nonlinear relationships between the technical indicators and the trading position. The cost is a substantial increase in the number of parameters in the model. For example, with $n = 5$ nodes in the network, and all 10 inputs, there will be $(5+1) + 5 \times (10+1) = 61$ coefficients in the model. This poses an obvious risk of overfitting, and hence erratic out-of-sample model performance. The parameters of the neural network are estimated iteratively using a least-squares criterion.

Third, since the risk in the model trades arises from the length of time the positions are held, we could use knowledge of analyst behaviour to limit this time. Specifically, we use the models to open long or short positions, but then close the positions when a profit target or stop loss level is hit. We have taken these to be the average limits applied by the analysts in the pre-test period (hours 1-1000). For the Schatz, the profit

target is on average 11bps away from the entry price, and the stop-loss level on average 7bps away. For the more volatile Bobl and Bund the figures are (21bps, 15bps) and (32bps, 20bps) respectively.

Table 10 summarises the results of these experiments. Taking positions only when analyst and model agree results in profits that are similar to, and sometimes lower than, analyst profits, with comparable levels of risk (drawdown). Adding the lagged dependent variable to the ordered response model naturally improves the statistical fit, but does not improve profit performance. In most markets and time periods profits are lower than, and drawdown is about the same as, those for the simple model. Neither of these procedures can be considered an improvement on the simple ordered choice model.

Table 10. Return and risk under alternative judgmental bootstrap models

		Schatz		Bobl		Bund	
		<i>Profit</i>	<i>Drawdown</i>	<i>Profit</i>	<i>Drawdown</i>	<i>Profit</i>	<i>Drawdown</i>
Analyst	<i>In-sample</i>	221	-26	220	-133	14	-237
	<i>Out-of-sample</i>	68	-45	238	-62	438	-75
Model	<i>In-sample</i>	277	-99	421	-170	662	-248
	<i>Out-of-sample</i>	112	-108	545	-250	532	-408
Model+Analyst	<i>In-sample</i>	157	-28	136	-89	113	-247
	<i>Out-of-sample</i>	85	-35	115	-39	204	-61
Model+ Lagged Analyst	<i>In-sample</i>	153	-127	339	-299	455	-370
	<i>Out-of-sample</i>	180	-98	274	-135	685	-314
Model Pooled across Analysts	<i>In-sample</i>	237	-148	421	-172	781	-250
	<i>Out-of-sample</i>	35	-178	510	-221	585	-239
Neural Network	<i>In-sample</i>	108	-22	254	-147	402	-211
	<i>Out-of-sample</i>	234	-64	108	-48	355	-391

Notes: The table shows levels of in-sample and out-of-sample profits and maximum drawdown in basis points. The Model+Analyst takes a position only if the restricted judgmental bootstrap model and the analyst agree. The Model+Lagged Analyst uses a judgmental bootstrap model with the lagged analyst position as a regressor. The Pooled Model uses data from all three markets (and hence the two analysts FAB and SKY) to estimate a common restricted ordered response model.

Pooling the data from both analysts is more productive. The best restricted model is simple and credible, containing only the current price: moving average ratio, the current stochastic, and the lagged short:long moving average ratio. The model's performance on the Schatz and Bobl (followed by analyst FAB) is only slightly inferior to the unpooled model. But for the Bund (followed by analyst SKY) not only are profits in- and out-of-sample higher for the pooled model, but the level of drawdown is also substantially reduced. So the bootstrap process has allowed superior expertise to be transferred from one analyst to another.

The last row of Table 10 show results from the single layer neural network. The number of nodes is chosen to minimise the prediction error for a test set of 15% of observations within the in-sample range of hours 1001-4000. After the overall architecture of the model is established, variables are progressively deleted to obtain a nonlinear restricted model. The final models are relatively simple, with 2-4 nodes, and 4-6 inputs. Even so, their predictive performance is erratic, and it does not appear that our judgmental bootstrap models can be improved by using a highly nonlinear functional form.

Finally, imposing typical stop loss and limit levels, and closing out positions at the end of each trading session, makes the model-based trading profile more realistic. However, shortening the periods when the model is in the market reduces profits sharply relative to analyst and unconstrained model results, and cannot be regarded as an improvement over, say, the pooled model.

6.6 CONCLUSIONS

This study has attempted to apply simple models of analyst decisions to the technical trading recommendations of two analysts. The environment of these decisions is rather more complex and ill-defined than in previous "clinical" bootstrapping exercises. In particular, although we have general evidence on the range of tools used

by the analysts, we do not know what indicator supported each decision to buy or sell. This contrasts with the clinical case where doctors can generally articulate a finite list of symptoms to support their diagnosis. Even in the canonical bankruptcy prediction problem in finance, the bankruptcy event is generally predictable on the basis of 5-10 financial ratios.

Perhaps not surprisingly we have failed in important respects the decisions of the analysts. In particular the bootstrap models we have developed spend much longer in open positions, and so are more exposed to market risk than the analysts. Although we have developed models that appear to make more money both in and out of sample, it is quite unclear whether on a risk-adjusted basis we have effected any improvement.

CHAPTER 7

SUMMARY

This thesis has tested three hypotheses about the performance of technical analysts, using directly observed data on their beliefs, their forecasts, and their trading recommendations. The hypotheses relate to the accuracy and rationality of forecasts (Chapter 4), the relationship between forecasts and trading recommendations (Chapter 5, and the relationship between trading recommendations and conventional mechanical indicators (Chapter 6). A typical analyst uses several indicators in preparing a forecast. These differ across analysts and in some cases change over time depending on market conditions. Analyst forecasts look little better than random, and do not appear technically rational. However, the analysts do possess market timing ability that lets them earn excess risk adjusted profits (Chapter 4). It seems difficult to adequately explain or model this ability. Simple rules based on lagged technical indicators do not fit analyst trading recommendations at all well (Chapter 6). Even rules based on the analysts own statements of where critical support and resistance lines lie do not outperform analyst judgment (Chapter 5). So one major conclusion from this study is that whatever value technical analysts have is difficult to capture using simple (and in the case of neural networks, quite complicated) statistical models.

The popularity of technical analysis is well documented in previous surveys as seen in Taylor and Allen (1992) and Lui and Mole (1998). In our survey results, there are a few interesting points worth noted. The technical methods used by our panel are mostly concentrated on the traditional tools suggested in the 1950s text look like Edward and Megee. There is very little evidence of diverting to the more “esoteric” approach like Elliott Wave and Gann analysis. However, the candlestick methods popularized in the 1990s have gained attention. Part of the reason could due to the accessibility/availability in various charting machines or a lot of the trading

recommendations are borne out by analyst experience and market timing rather than the sophistication of the methods themselves. The positive result of predicted turning point cited by Lui and Mole is however not being totally agreed by our study. The lack of calibration in range forecast could be due to the extrapolative nature as mentioned in previous research. In terms of the forecasting performance from our panel, this is no firm signal of domination by a single analyst or a group of analysts and the performance rating is diverse, this is in line with Allen and Taylor (1990).

Overall, however, our findings cast doubt on the value of many empirical studies of technical trading rules. Following the trading recommendations of the group of analysts studied here would have yielded consistent profits, after allowing for slippage and transactions costs. In contrast, none of the synthetic rules we examined yields significant profits. It would be naïve to think that in a speculatively efficient market all technical analysts could be as successful as those in our sample. But we should recognise the danger that a mechanical interpretation of technical indicators may systematically understate their value, relative to what can be achieved by professional analysts combining technical indicators with other elements of judgement.

None of the mechanical rules looks “real” judging from the support and resistance levels based on the morning comments. To varying degrees, they fail to capture key characteristics of the analysts’ recommended trading positions. Filter rules and directional trading strategies involve excessively long holding periods, and the profits across markets from these rules bear no relation to profits achieved by analysts. Support and resistance breaking rules oversimplify the way that these levels are used in practice, and again fail to reflect the level and pattern of potential profits from trading on technical indicators. One reason could be that these levels are used as an indication of the market trend rather than entry and exit points for trading. We simply cannot use a single set of rules to generate profits in this highly speculative market.

As for the bond paper, we look at trading ideas generated from analysts to compare with statistical models. The accuracy of direction is again statistically insignificant

despite better profit produced than the naïve buy/hold strategy but the sequence of the recommendations is able withstand the bootstrapping shuffling process which unable to duplicate the profits made from analysts positions.

The ordered response model uses some of the mechanical trading rules mentioned in the survey in chapter 4. The purpose is not to mimic how the decisions are made rather to mimic analysts' decision using a small number of quantifiable inputs. It seems that simple models parameterised on analysts trading recommendations can produce at least as much profit as the analysts themselves, and in most cases substantially more. However, the pattern of model-based trades suggests that this excess profit is not achieved without an increase in risk. The final level of (restricted) model profits is above that of the analyst in both in-sample and out-of-sample periods in all cases. But because the model is in the market continuously, the volatility of profits is high, and quite large losses can be experienced *en route* to the final higher level of profits.

There are a few suggestions that could be learned from the thesis. The survey results can be done on a bigger scale than the one that I did and the participants could be more diverse than the sample despite our penal of analysts are very experienced in this field. We managed to capture 44 weeks of survey results with 14 analysts giving more than 20 weeks of forecast information, this is similar to the time length in previous Taylor and Allen survey. It would be interesting to see a longer term survey in the future.

There are also several unexploited features of the survey. For example, we have investigated but not pursued the question of whether analysts are well calibrated in the sense that their high/low ranges change appropriately from week to week as volatility changes. Nor have we examined the value of the support and resistance levels, or market tone, reported in the survey, partly because we had access to the much larger data set on support and resistance studies in Chapter 5.

The findings of the currency paper suggest some more fruitful avenues of research on technical analysis. The relationship between the analysts' trading position and support and resistance levels is complex but does have some logic. Sometimes trades are undertaken after support or resistance is broken, with the implication that prices are expected to continue their trend. Sometimes, trades are undertaken when support and resistance levels are not broken, with the implication that prices are expected to reverse their local trend. How analysts can distinguish one situation from another, and more generally whether their behaviour is consistent enough to be captured by an expert system, are interesting questions for research. Another issue concerns the role of limit orders. Much attention is given in academic studies to directional forecasting and decisions over position-taking, but little attention is paid to the role of position management, and the determination of stop loss and limit levels. All of our analysts' trades are circumscribed by tight stop loss and limit orders. It is interesting to ask how they are determined in practice, and they might be ideally determined. Moreover, we have not looked at the wording of the morning commentaries attached with the technical signals in detail. A systematic textual analysis might give us better insight of the "reason" behind the trading decision as compared to the simple technical breakout rules investigated here.

Similarly, in modelling the recommendations in the bond futures markets we have mainly used simple linear models. Our data would be amenable to modelling using a more sophisticated expert system, and given the failure of the simple models, this would be a natural next step in our research.

REFERENCES

- Acar, E. 1998, Expected returns of directional forecasters, in E. Acar and S. Satchell (Eds.), *Advanced Trading Rules*, Oxford: Butterworth-Heinemann, 51-80.
- Aitchison, J. and S. D. Silvey, 1957, The generalization of probit analysis to the case of multiple responses, *Biometrika*, 44, 131-40.
- Alexander, S. S., 1961, Price movements in speculative markets: trends or random walks, *Industrial Management Review*, 2, 7-26.
- Allen F. and R. Karjaleinen, 1999, Using genetic algorithms to find technical trading rules, *Journal of Financial Economics*, 51, 2, 245-271.
- Allen H. and Taylor M.P., 1990, Charts, Noise and Fundamentals in the London Foreign Exchange Market, *Economic Journal*, 49-59.
- Armstrong, J. S., 2001, Judgmental bootstrapping: inferring experts' rules for forecasting, in J. S. Armstrong (Ed), *Principles of Forecasting: a Handbook for Researchers and Practitioners*, Kluwer: Dordrecht.
- Batchelor, R. and Dua, P., 1992, Blue chip rationality tests, *Journal of Money, Credit and Banking*, Vol.23, pages 692-705.
- Batchelor, R. A., 1990, All forecasters are equal, *Journal of Business and Economic Statistics*, 8, 1, 143-4.
- Batchelor, R. A. and T. Y. Kwan, 2001, *A survey of technical analysts*, Cass Business School Discussion Paper 01-15.
- Bensaid, B. and O. De Bandt, 1998, Stop loss rules as a monitoring device: theory and evidence from the bond futures market, in E. Acar and S. Satchell (Eds.), *Advanced Trading Rules*, Oxford: Butterworth-Heinemann.
- Boothe, P., 1983, Speculative profit opportunities in the Canadian foreign exchange market, *Canadian Journal of Economics*, 16, 4, 603-11.
- Boothe P. and D. Glassman, 1987, Comparing exchange rate forecasting models: accuracy versus profitability, *International Journal of Forecasting*, 3, 1, 295-320.
- Brock W. A., LeBaron B. and J. Lakonishok, 1992, Simple technical trading rules and the stochastic properties of stock returns, *Journal of Finance*, 47, 1731-1764.

Brown S.J., Goetzmann.W.N. and Kumar A.,1998,The Dow Theory: William Peter Hamilton's Track Record Reconsidered, *Journal of Finance*, vol.LIII, no.4, 1311 – 1333.

Chang, P. H. K. and C. L. Osler, 1999, Methodical madness: technical analysis and the irrationality of exchange rate forecasts, *Economic Journal*, 109, 636-661.

Cheung, Y-W. and M. D. Chinn, 2001, Currency traders and exchange rate dynamics: a survey of the US market, *Journal of International Money and Finance*, 20, 439-471.

Cheung, Y-W, M. D. Chinn and I. Marsh, 1999, *How do UK-based foreign exchange dealers think their market operates?*, National Bureau of Economic Research Working Paper 7524.

Chinn, Menzie and Frankel, Jeffrey, 1991, Are Exchange Rate Expectations Biased? Tests for a Cross-Sections of 25 Currencies, *NBER Working Paper* no.3807.

Cowles A.,1944, Stock market forecasting, *Econometrica*,12,206-214

Cowles A.,1933, Can stock market forecasters forecast? ,*Econometrica*,1,309-324

Curcio, R. and C. Goodhart, 1992, *When support/ resistance levels are broken can profits be made? Evidence from the foreign exchange market*, LSE Financial Markets Group Discussion Paper 142.

Curcio, R., C. Goodhart, D. Guillaume and R. Payne, 1997, Do technical trading rules generate profits? Conclusions from the intraday foreign exchange market, *International Journal of Financial Economics*, 2, 267-280.

Dawes, R. M., 1971, A case study of graduate admissions: applications of three principles of human decision-making, *American Psychologist*, 26, 180-88.

Dawes, R. M. and B. Corrigan, 1974, Linear models in decision-making, *Psychological Bulletin*, 81, 95-106.

de Grauwe, P. and D. Decupere, 1992, Psychological barriers in the foreign exchange markets, *Journal of International and Comparative Economics*, 1, 87-101.

Donaldson, R. G. and H. Y. Kim, 1993, Price barriers in the Dow Jones Industrial Average, *Journal of Financial and Quantitative Analysis*, 28, 3, 313-30.

Edwards, R. and J. Magee, 2001, *Technical analysis of stock trends*, Amacon (8th Edition).

- Fama, E. and M. E. Blume, 1966, Filter rules and stock market trading, *Journal of Business*, 39, 226-41.
- Frankel, J.A. and Froot, K.A., 1990, The rationality of the foreign exchange rate: Chartists, Fundamentalists, and Trading in the Foreign Exchange Market, *American Review Economic*, 181-185.
- Gencay, R., 1996, Nonlinear prediction of security returns with moving average rules, *Journal of Forecasting*, 15, 165-174.
- Gencay, R., 1998, The predictability of security returns with simple technical trading rules, *Journal of Empirical Finance*, 5(4), 347-59.
- Gehrig, T. and Menkhoff, L. 2003, *Technical Analysis in Foreign Exchange –The Workhorse Gains Further Ground*, Universitat Freiburg und CEPR, London and Universitat Hannover, Discussion Paper No. 278.
- Grove, W. M., and P. E. Meehl, 1996, Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: the clinincal-statistical controversy, *Psychology, Public Policy and Law*, 2, 293-323.
- Grove, W. M., D. H. Zald, B. S. Lebow, B. E. Snits and C. E. Nelson, 2000, Clinical vs. mechanical prediction: a meta-analysis, *Psychological Assessment*, 12, 19-30.
- Gunasekarage, A. and D. Power, 2001, The profitability of moving average trading rules in South Asian stock markets, *Emerging Markets Review*, 2(1), 17-33.
- Hafer, R. W., S. E. Hein and S. S. MacDonald, 1992, 'Market and survey forecasts of the three-month Treasury bill rate', *Journal of Business*, 65, 1, 123-138.
- James, F. E., 1968, Monthly moving averages - an effective investment tool?, *Journal of Financial and Quantitative Analysis*, 3, 3, 315-326.
- Hafer, R. W., S. E. Hein and S. S. MacDonald, 1992, 'Market and survey forecasts of the three-month Treasury bill rate', *Journal of Business*, 65, 1, 123-138.
- Karpoff, J. M., 1987, The relation between price changes and trading volume: a survey, *Journal of Financial and Quantitative Analysis*, 22, 109-126.
- Keane, M.P. and Runkle, D., 1990, Testing the rationality of price forecasts: new evidence from panel data, *American Economic Review*, Vol.80, pages 714-35.
- Kolb, R. A. and H. O. Stekler, 1996, 'How well do analysts forecast interest rates?', *Journal of Forecasting*, 15, 385-94.

- LeBaron, B., 1999, Technical trading rule profitability of foreign exchange intervention, *Journal of International Economics*, 49, 125-43.
- Leitch, G. and J. E. Tanner, 1991, Economic forecast evaluation: profit versus conventional error measures, *American Economic Review*, 81, 580-590.
- Levich, R. M. and L. R. Thomas, 1993, The significance of technical trading rule profits in the foreign exchange market; a bootstrap approach, *Journal of International Money and Finance*, 12, 451-474.
- Lee, C. I. and I. Mathur, 1994, Trading rule profits in European currency spot cross rates, *Journal of Banking and Finance*, 20(5), 949-62.
- Lee, C. I., K. C. Gleason and I. Mathur, 2001, Trading rule profits in Latin American currency markets, *International Review of Financial Analysis*, 10(2), 135-156.
- Lee, C. I., M-S. Pau, and Y. A. Liu, 2001, On market efficiency of Asian foreign exchange rates: evidence from a joint variance ratio test and technical trading rules, *Journal of International Financial Markets, Institutions and Money*, 11(2) 199-214.
- Levy, R.A. (1971). The predictive Significance of Five-Point Chart Patterns, *Journal of Business*, 316-323.
- Liu, Y-H. and D. Mole, 1998, The use of fundamental and technical analysis by foreign exchange dealers: Hong Kong evidence, *Journal of International Money and Finance*, 17, 535-545.
- Lo, A., H. Mamaysky and Jiang Wang, 2000, Foundations of technical analysis: computational algorithms, statistical inference and empirical implementation, *Journal of Finance*, 55, 1705-1765.
- Louis P.Lukac, B.Wade Brorsen and Scott H. Irwin, 1988, Similarity of computer guided technical trading systems, *The journal of futures markets*, 8, 1-18.
- Lyons, R. K., 1998, Profits and position control: a week of FX dealing, *Journal of International Money and Finance*, 17, 1, 97-115.
- MacDonald, S. S. and Hein, S. E., 1989, 'Futures rates and forward rates as predictors of near-term Treasury bill rates', *Journal of Futures Markets*, 9, 249-62.
- Mandelbrot, Benoit, 1966, Forecast of future prices, Unbiased markets and 'Martingale' models, *Journal of Business*, vol. XXXIX, no.1, Part II (Supplement).

- Malkiel, Burton G., 1996, *A random walk down wall street*, W.W.Norton & Company, Inc.
- Meehl, P. E., 1954, *Clinical versus statistical prediction: a theoretical analysis and a review of the evidence*, republished (1996) by Jason Aronson: Northvale New Jersey.
- Menkhoff, L., 1997, Examining the use of technical currency analysis, *International Journal of Finance and Economics*, 2, 4, 307-18.
- Murphy, J.A.,1986, Futures Fund Performance: A Test of the Effectiveness of Technical Analysis, *The Journal of Futures Markets*, vol 6, no. 2, 175-185.
- Murphy, J.J.,1986, *Technical Analysis of the futures markets*, New York Institute of Finance.
- Murphy, J., 1998, *Technical analysis of the financial markets*, New York Institute of Finance, (2nd edition).
- Neely, C. J., 1997, Technical analysis in the foreign exchange market: a layman's guide, *Federal Reserve Bank of St. Louis Review*, Sep-Oct, 23-38.
- Neely, C. J., P. Weller and R. Ditmar, 1997, Is technical analysis in the foreign exchange market profitable? A genetic programming approach, *Journal of Financial and Quantitative Analysis*, 32, 4, 405-26.
- Neely, C. J. and P. Weller, 1999, Technical trading rules in the European Monetary System, *Journal of International Money and Finance*, 18, 3, 429-58.
- Neely,C. J. and P. Weller, 1999, Technical Analysis and Central Bank Intervention, Federal Reserve Bank of St.Louis Working Paper-97-002B.
- Nison, S., 2001, *Japanese Candlestick Charting Techniques*, Prentice Hall.
- Oberlechner, T., 2001, Importance of technical and fundamental analysis in the European exchange market, *International Journal of Finance and Economics*, 6, 1, 81-93.
- Olson, D., 2003, Have trading rule profits in the currency market declined over time?, *Journal of Banking and Finance*, 28,85-105.
- Osler, C.L. and P.H.Kevin Chang,1995, *Head and Shoulders: Not just a flaky pattern*, Federal Reserve Bank of New York.
- Osler, C.L.,1998, *Identifying Noise Traders: The Head an Shoulders Pattern in U.S. equities*, Federal Reserve Bank of New York.

Osler, C., 2000, Support for Resistance: Technical Analysis and Intraday Exchange Rates, *Federal Reserve Bank of New York Economic Policy Review*, July, Volume 6, Number 2, 53-69.

Osler, C. L., 2001, *Currency Orders and Exchange-Rate Dynamics: Explaining the Success of Technical Analysis*, Federal Reserve bank of New York Staff Report No. 125, April 2001.

Prechter, R., and A. J. Frost, 2000, *Elliott Wave Principle: the Key to Market Behaviour*, Wiley: New York, 10th Ed.

Pring, M., 1998, *Introduction to technical analysis*, McGraw Hill Professional.

Pruitt S.W. and White R.E., 1988, The CRISMA trading system: who says technical analysis can't bear the market? *The Journal of Portfolio Management*, 55-58.

Silber W.L., 1994, Technical Trading: When it works and When it doesn't, *The Journal of Derivatives*, 39-44.

Sullivan, R., A. Timmerman and H. White, 1999, Data-snooping, technical trading rule performance and the bootstrap, *Journal of Finance*, 54, 5.

Sweeney, R. J., 1986, Beating the foreign exchange market, *Journal of Finance*, 41, 163-82.

Sweeney, R. J., 1988, Some new filter rule tests: method and results, *Journal of Financial and Quantitative Analysis*, 23, 285-300.

Szakmary, A. C. and I. Mathur, 1997, Central bank intervention and trading rule profits in foreign exchange markets, *Journal of International Money and Finance*, 16, 4, 515-535.

Taylor, M. P. and H. Allen H, 1992, The use of technical analysis in the foreign exchange market, *Journal of International Money and Finance*, 11, 3, 304-314.

Taylor S., 1994, Trading futures using a channel rule: a study of the predictive power of technical analysis with currency examples, *Journal of Futures Markets*, 14, 215-235.